Red Hat Enterprise Virtualization 3.3 Installation Guide

Installing Red Hat Enterprise Virtualization Environments

Andrew Burden
Steve Gordon
Tim Hildred
Cheryn Tan

Red Hat Enterprise Virtualization 3.3 Installation Guide

Installing Red Hat Enterprise Virtualization Environments

Andrew Burden
aburden@redhat.com

Steve Gordon
sgordon@redhat.com

Tim Hildred
thildred@redhat.com

Cheryn Tan
chetan@redhat.com

Legal Notice

Copyright © 2014 Red Hat, Inc.

This document is licensed by Red Hat under the Creative Commons Attribution-ShareAlike 3.0 Unported License. If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, MetaMatrix, Fedora, the Infinity Logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL ® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack Logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

A comprehensive guide to installing Red Hat Enterprise Virtualization environments.

Table of Contents

Preface **7**
1. Document Conventions 7
1.1. Typographic Conventions 7
1.2. Pull-quote Conventions 8
1.3. Notes and Warnings 9
2. Getting Help and Giving Feedback 9
2.1. Do You Need Help? 9
2.2. We Need Feedback! 10

Part I. Before you Begin **11**

Chapter 1. Introduction **12**
1.1. Red Hat Enterprise Virtualization Architecture 12
1.2. Red Hat Enterprise Virtualization System Components 13
1.3. Red Hat Enterprise Virtualization Resources 13
1.4. Red Hat Enterprise Virtualization API Support Statement 14
1.5. Introduction to Virtual Machines 15
1.6. Supported Virtual Machine Operating Systems 16
1.7. Red Hat Enterprise Virtualization Installation Workflow 17

Chapter 2. System Requirements **19**
2.1. Workflow Progress — System Requirements 19
2.2. Hardware Requirements 19
2.2.1. Red Hat Enterprise Virtualization Hardware Requirements Overview 19
2.2.2. Red Hat Enterprise Virtualization Manager Hardware Requirements 20
2.2.3. Virtualization Host Hardware Requirements 20
2.2.3.1. Virtualization Host Hardware Requirements Overview 20
2.2.3.2. Virtualization Host CPU Requirements 21
2.2.3.3. Virtualization Host RAM Requirements 22
2.2.3.4. Virtualization Host Storage Requirements 22
2.2.3.5. Virtualization Host PCI Device Requirements 23
2.3. Software Requirements 24
2.3.1. Red Hat Enterprise Virtualization Operating System Requirements 24
2.3.2. Red Hat Enterprise Virtualization Manager Client Requirements 24
2.3.3. Red Hat Enterprise Virtualization Manager Software Channels 25
2.3.4. Directory Services 26
2.3.4.1. About Directory Services 26
2.3.4.2. Directory Services Support in Red Hat Enterprise Virtualization 26
2.3.5. Firewall Configuration 29
2.3.5.1. Red Hat Enterprise Virtualization Manager Firewall Requirements 29
2.3.5.2. Virtualization Host Firewall Requirements 30
2.3.5.3. Directory Server Firewall Requirements 32
2.3.5.4. Database Server Firewall Requirements 33
2.3.6. Required User Accounts and Groups 33
2.3.6.1. Red Hat Enterprise Virtualization Manager User Accounts 33
2.3.6.2. Red Hat Enterprise Virtualization Manager Groups 33
2.3.6.3. Virtualization Host User Accounts 34
2.3.6.4. Virtualization Host Groups 34

Part II. Installing Red Hat Enterprise Virtualization Manager **36**

Chapter 3. Manager Installation **37**
3.1. Workflow Progress — Installing Red Hat Enterprise Virtualization Manager 37

3.2. Installing the Red Hat Enterprise Virtualization Manager 37
3.3. Subscribing to the Red Hat Enterprise Virtualization Channels 40
3.3.1. Subscribing to the Red Hat Enterprise Virtualization Manager Channels using Subscription Manager 40
3.3.2. Subscribing to the Red Hat Enterprise Virtualization Manager Channels Using RHN Classic 41
3.4. Installing the Red Hat Enterprise Virtualization Manager Packages 44
3.5. Configuring Red Hat Enterprise Virtualization Manager 45
3.6. Passwords in Red Hat Enterprise Virtualization Manager 49
3.7. Preparing a PostgreSQL Database for Use with Red Hat Enterprise Virtualization Manager 50
3.8. Configuring the Manager to Use a Manually Configured Local or Remote PostgreSQL Database 51
3.9. Connecting to the Administration Portal 52
3.10. Removing Red Hat Enterprise Virtualization Manager 53

Chapter 4. Self-Hosted Engine 55
4.1. About the Self-Hosted Engine 55
4.2. Limitations of the Self-Hosted Engine 55
4.3. Installing the Self-Hosted Engine 55
4.4. Configuring the Self-Hosted Engine 56
4.5. Migrating to a Self-Hosted Environment 62
4.6. Installing Additional Hosts to a Self-Hosted Environment 71
4.7. Maintaining the Self-Hosted Engine 72

Chapter 5. Data Collection Setup and Reports Installation 74
5.1. Workflow Progress — Data Collection Setup and Reports Installation 74
5.2. Data Collection Setup and Reports Installation Overview 74
5.3. Installing and Configuring the History Database 75
5.4. Installing and Configuring Red Hat Enterprise Virtualization Manager Reports 76

Chapter 6. Updating the Red Hat Enterprise Virtualization Environment 78
6.1. Upgrades between Minor Releases 78
6.1.1. Checking for Red Hat Enterprise Virtualization Manager Updates 78
6.1.2. Updating Red Hat Enterprise Virtualization Manager 78
6.1.3. Troubleshooting for Upgrading Red Hat Enterprise Virtualization Manager 80
6.1.4. Updating Red Hat Enterprise Virtualization Manager Reports 80
6.1.5. Updating Red Hat Enterprise Virtualization Hypervisors 81
6.1.6. Updating Red Hat Enterprise Linux Virtualization Hosts 83
6.1.7. Updating the Red Hat Enterprise Virtualization Guest Tools 83
6.2. Upgrading to Red Hat Enterprise Virtualization 3.3 84
6.2.1. Red Hat Enterprise Virtualization Manager 3.3 Upgrade Overview 84
6.2.2. Red Hat Enterprise Virtualization 3.3 Upgrade Considerations 84
6.2.3. Upgrading to Red Hat Enterprise Virtualization Manager 3.3 85
6.3. Upgrading to Red Hat Enterprise Virtualization Manager 3.2 87
6.3.1. Upgrading to Red Hat Enterprise Virtualization Manager 3.2 87
6.4. Upgrading to Red Hat Enterprise Virtualization Manager 3.1 89
6.4.1. Upgrading to Red Hat Enterprise Virtualization Manager 3.1 89
6.5. Post-upgrade Tasks 93
6.5.1. Features Requiring a Compatibility Upgrade to Red Hat Enterprise Virtualization 3.3 93
6.5.2. Changing the Cluster Compatibility Version 94
6.5.3. Changing the Data Center Compatibility Version 95

Part III. Installing Virtualization Hosts 97

Chapter 7. Introduction to Virtualization Hosts 98

7.1. Workflow Progress — Installing Virtualization Hosts 98
7.2. Introduction to Virtualization Hosts 98

Chapter 8. Installing Red Hat Enterprise Virtualization Hypervisor Hosts 100
8.1. Red Hat Enterprise Virtualization Hypervisor Installation Overview 100
8.2. Installing the Red Hat Enterprise Virtualization Hypervisor Packages 100
8.3. Preparing Hypervisor Installation Media 102
8.3.1. Preparing USB Hypervisor Installation Media 102
8.3.1.1. Preparing a Hypervisor USB Storage Device 102
8.3.1.2. Preparing USB Installation Media Using livecd-iso-to-disk 102
8.3.1.3. Preparing USB Installation Media Using dd 104
8.3.1.4. Preparing USB Installation Media Using dd on Linux Systems 105
8.3.1.5. Preparing USB Installation Media Using dd on Windows Systems 106
8.3.2. Preparing Optical Hypervisor Installation Media 108
8.3.3. Booting from Hypervisor Installation Media 109
8.3.3.1. Booting the Hypervisor from USB Installation Media 109
8.3.3.2. Booting the Hypervisor from Optical Installation Media 109
8.3.3.3. Troubleshooting BIOS Settings and Boot Process 110
8.3.3.4. Choosing Hypervisor Boot Options 111
8.4. Installing the Hypervisor 114
8.4.1. Hypervisor Menu Actions 114
8.4.2. Installing the Hypervisor 115
8.5. Configuring the Hypervisor 118
8.5.1. Logging into the Hypervisor 118
8.5.2. Selecting Hypervisor Keyboard 118
8.5.3. Viewing Hypervisor Status 119
8.5.4. Configuring Hypervisor Network 119
8.5.4.1. Hypervisor Network Screen 120
8.5.4.2. Configuring Hypervisor Host Name 120
8.5.4.3. Configuring Hypervisor Domain Name Servers 120
8.5.4.4. Configuring Hypervisor Network Time Protocol 121
8.5.4.5. Configuring Hypervisor Network Interfaces 121
8.5.5. Configuring Hypervisor Security 123
8.5.6. Configuring Hypervisor Simple Network Management Protocol 123
8.5.7. Configuring Hypervisor Common Information Model 124
8.5.8. Configuring Logging 125
8.5.9. Configuring the Hypervisor for Red Hat Network 126
8.5.10. Configuring Hypervisor Kernel Dumps 128
8.5.11. Configuring Hypervisor Remote Storage 128
8.6. Attaching the Hypervisor to the Red Hat Enterprise Virtualization Manager 129
8.6.1. Configuring Hypervisor Management Server 129
8.6.2. Using the Hypervisor 130
8.6.3. Approving a Hypervisor 130

Chapter 9. Installing Red Hat Enterprise Linux Hosts 132
9.1. Red Hat Enterprise Linux Hosts 132
9.2. Host Compatibility Matrix 132
9.3. Preparing a Red Hat Enterprise Linux Host 132
9.3.1. Installing Red Hat Enterprise Linux 132
9.3.2. Subscribing to Required Channels Using Subscription Manager 134
9.3.3. Subscribing to Required Channels Using RHN Classic 136
9.3.4. Configuring Virtualization Host Firewall 138
9.3.5. Configuring Virtualization Host sudo 140
9.3.6. Configuring Virtualization Host SSH 141

9.3.6. Configuring Virtualization Host SSH 141
9.4. Adding a Red Hat Enterprise Linux Host 142
9.5. Explanation of Settings and Controls in the New Host and Edit Host Windows 143
9.5.1. Host General Settings Explained 143
9.5.2. Host Power Management Settings Explained 144
9.5.3. SPM Priority Settings Explained 146
9.5.4. Host Console Settings Explained 147

Part IV. Environment Configuration 148

Chapter 10. Planning your Data Center 149
10.1. Workflow Progress — Planning Your Data Center 149
10.2. Planning Your Data Center 149
10.3. Data Centers 150
10.3.1. Data Centers in Red Hat Enterprise Virtualization 150
10.3.2. Creating a New Data Center 151
10.4. Clusters 151
10.4.1. Clusters in Red Hat Enterprise Virtualization 151
10.4.2. Creating a New Cluster 152
10.4.3. Enabling Gluster Processes on Red Hat Storage Nodes 153

Chapter 11. Network Setup 155
11.1. Workflow Progress — Network Setup 155
11.2. Networking in Red Hat Enterprise Virtualization 155
11.3. Logical Networks 157
11.3.1. Creating a New Logical Network in a Data Center or Cluster 157
11.3.2. Editing Host Network Interfaces and Adding Logical Networks to Hosts 158
11.3.3. Explanation of Settings and Controls in the General Tab of the New Logical Network and Edit Logical Network Windows 160
11.3.4. Editing a Logical Network 161
11.3.5. Designate a Specific Traffic Type for a Logical Network with the Manage Networks Window
11.3.6. Explanation of Settings in the Manage Networks Window 163 161
11.3.7. Adding Multiple VLANs to a Single Network Interface Using Logical Networks 163
11.3.8. Multiple Gateways 165
11.4. Using the Networks Tab 165
11.4.1. Importing Networks from External Providers 166
11.4.2. Limitations to Importing Networks from External Providers 166
11.5. Bonds 167
11.5.1. Bonding Logic in Red Hat Enterprise Virtualization 167
11.5.2. Bonding Modes 168
11.5.3. Creating a Bond Device Using the Administration Portal 169
11.5.4. Example Uses of Custom Bonding Options with Host Interfaces 170

Chapter 12. Storage Setup 172
12.1. Workflow Progress — Storage Setup 172
12.2. Introduction to Storage in Red Hat Enterprise Virtualization 172
12.3. Adding Storage to the Environment 174
12.3.1. Adding NFS Storage 174
12.3.1.1. Preparing NFS Storage 174
12.3.1.2. Attaching NFS Storage 176
12.3.2. Adding pNFS Storage 177
12.3.2.1. Preparing pNFS Storage 177
12.3.2.2. Attaching pNFS Storage 178
12.3.3. Adding iSCSI Storage 179
12.3.4. Adding FCP Storage 181

12.3.5. Adding Local Storage 183
12.3.5.1. Preparing Local Storage 183
12.3.5.2. Adding Local Storage 184
12.3.6. Adding POSIX Compliant File System Storage 185
12.3.6.1. POSIX Compliant File System Storage in Red Hat Enterprise Virtualization 185
12.3.6.2. Attaching POSIX Compliant File System Storage 186
12.4. Populating the ISO Domain 187
12.4.1. Populating the ISO Storage Domain 187
12.4.2. VirtIO and Guest Tool Image Files 188
12.4.3. Uploading the VirtIO and Guest Tool Image Files to an ISO Storage Domain 188

Log Files 190
A.1. Red Hat Enterprise Virtualization Manager Installation Log Files 190
A.2. Red Hat Enterprise Virtualization Manager Log Files 190
A.3. Red Hat Enterprise Virtualization Host Log Files 191
A.4. Remotely Logging Host Activities 191
A.4.1. Setting Up a Virtualization Host Logging Server 191
A.4.2. Configuring Logging 192
A.4.3. Configuring Logging 193

Additional Utilities 195
B.1. Domain Management Tool 195
B.1.1. What is the Domain Management Tool? 195
B.1.2. Syntax for the Domain Management Tool 195
B.1.3. Adding Domains to Configuration 197
B.1.4. Editing a Domain in the Configuration 198
B.1.5. Deleting a Domain from the Configuration 198
B.1.6. Validating Domain Configuration 199
B.1.7. Listing Domains in Configuration 200
B.2. Configuration Tool 200
B.2.1. Configuration Tool 200
B.2.2. Syntax for engine-config Command 201
B.3. Image Uploader 202
B.3.1. Virtual Machine Image Uploader 202
B.3.2. Syntax for the engine-image-uploader Command 202
B.3.3. Creating an OVF Archive That is Compatible with the Image Uploader 204
B.3.4. Basic engine-image-uploader Usage Examples 205
B.4. ISO Uploader 205
B.4.1. ISO Uploader 206
B.4.2. Syntax for engine-iso-uploader Command 206
B.4.3. Usage Examples 208
B.4.3.1. Specifying an NFS Server 208
B.4.3.2. Basic ISO Uploader Usage 208
B.5. Log Collector 209
B.5.1. Log Collector 209
B.5.2. Syntax for engine-log-collector Command 209
B.5.3. Basic Log Collector Usage 213
B.6. SPICE Proxy 213
B.6.1. SPICE Proxy Overview 213
B.6.2. SPICE Proxy Machine Setup 214
B.6.3. Turning on SPICE Proxy 214
B.6.4. Turning Off a SPICE Proxy 215
B.7. Squid Proxy 216
B.7.1. Installing and Configuring a Squid Proxy 216

B.7.1. Installing and Configuring a Squid Proxy 216

Revision History **221**

Preface

1. Document Conventions

This manual uses several conventions to highlight certain words and phrases and draw attention to specific pieces of information.

1.1. Typographic Conventions

Four typographic conventions are used to call attention to specific words and phrases. These conventions, and the circumstances they apply to, are as follows.

`Mono-spaced Bold`

Used to highlight system input, including shell commands, file names and paths. Also used to highlight keys and key combinations. For example:

> To see the contents of the file **`my_next_bestselling_novel`** in your current working directory, enter the **`cat my_next_bestselling_novel`** command at the shell prompt and press **`Enter`** to execute the command.

The above includes a file name, a shell command and a key, all presented in mono-spaced bold and all distinguishable thanks to context.

Key combinations can be distinguished from an individual key by the plus sign that connects each part of a key combination. For example:

> Press **`Enter`** to execute the command.
>
> Press **`Ctrl`**+**`Alt`**+**`F2`** to switch to a virtual terminal.

The first example highlights a particular key to press. The second example highlights a key combination: a set of three keys pressed simultaneously.

If source code is discussed, class names, methods, functions, variable names and returned values mentioned within a paragraph will be presented as above, in **`mono-spaced bold`**. For example:

> File-related classes include **`filesystem`** for file systems, **`file`** for files, and **`dir`** for directories. Each class has its own associated set of permissions.

Proportional Bold

This denotes words or phrases encountered on a system, including application names; dialog-box text; labeled buttons; check-box and radio-button labels; menu titles and submenu titles. For example:

> Choose **System** → **Preferences** → **Mouse** from the main menu bar to launch **Mouse Preferences**. In the **Buttons** tab, select the **`Left-handed mouse`** check box and click **`Close`** to switch the primary mouse button from the left to the right (making the mouse suitable for use in the left hand).
>
> To insert a special character into a **`gedit`** file, choose **Applications** → **Accessories** → **Character Map** from the main menu bar. Next, choose **Search** → **Find…** from the **Character Map** menu bar, type the name of the character in the **Search** field and click **`Next`**. The character you sought will be highlighted in the

> **Character Table**. Double-click this highlighted character to place it in the **Text to copy** field and then click the **Copy** button. Now switch back to your document and choose **Edit** → **Paste** from the **gedit** menu bar.

The above text includes application names; system-wide menu names and items; application-specific menu names; and buttons and text found within a GUI interface, all presented in proportional bold and all distinguishable by context.

Mono-spaced Bold Italic or ***Proportional Bold Italic***

Whether mono-spaced bold or proportional bold, the addition of italics indicates replaceable or variable text. Italics denotes text you do not input literally or displayed text that changes depending on circumstance. For example:

> To connect to a remote machine using ssh, type **ssh *username*@*domain.name*** at a shell prompt. If the remote machine is **example.com** and your username on that machine is john, type **ssh john@example.com**.
>
> The **mount -o remount *file-system*** command remounts the named file system. For example, to remount the **/home** file system, the command is **mount -o remount /home**.
>
> To see the version of a currently installed package, use the **rpm -q *package*** command. It will return a result as follows: ***package-version-release***.

Note the words in bold italics above: username, domain.name, file-system, package, version and release. Each word is a placeholder, either for text you enter when issuing a command or for text displayed by the system.

Aside from standard usage for presenting the title of a work, italics denotes the first use of a new and important term. For example:

> Publican is a *DocBook* publishing system.

1.2. Pull-quote Conventions

Terminal output and source code listings are set off visually from the surrounding text.

Output sent to a terminal is set in **mono-spaced roman** and presented thus:

```
books        Desktop   documentation  drafts  mss    photos   stuff  svn
books_tests  Desktop1  downloads      images  notes  scripts  svgs
```

Source-code listings are also set in **mono-spaced roman** but add syntax highlighting as follows:

```
static int kvm_vm_ioctl_deassign_device(struct kvm *kvm,
                  struct kvm_assigned_pci_dev *assigned_dev)
{
         int r = 0;
         struct kvm_assigned_dev_kernel *match;

         mutex_lock(&kvm->lock);

         match = kvm_find_assigned_dev(&kvm->arch.assigned_dev_head,
                                       assigned_dev->assigned_dev_id);
         if (!match) {
                 printk(KERN_INFO "%s: device hasn't been assigned
```

```
before, "
                      "so cannot be deassigned\n", __func__);
                    r = -EINVAL;
                    goto out;
            }

            kvm_deassign_device(kvm, match);

            kvm_free_assigned_device(kvm, match);

out:
            mutex_unlock(&kvm->lock);
            return r;
}
```

1.3. Notes and Warnings

Finally, we use three visual styles to draw attention to information that might otherwise be overlooked.

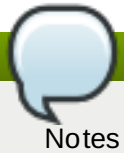

Note

Notes are tips, shortcuts or alternative approaches to the task at hand. Ignoring a note should have no negative consequences, but you might miss out on a trick that makes your life easier.

Important

Important boxes detail things that are easily missed: configuration changes that only apply to the current session, or services that need restarting before an update will apply. Ignoring a box labeled “Important” will not cause data loss but may cause irritation and frustration.

Warning

Warnings should not be ignored. Ignoring warnings will most likely cause data loss.

2. Getting Help and Giving Feedback

2.1. Do You Need Help?

If you experience difficulty with a procedure described in this documentation, visit the Red Hat Customer Portal at http://access.redhat.com. Through the customer portal, you can:

- search or browse through a knowledgebase of technical support articles about Red Hat products.
- submit a support case to Red Hat Global Support Services (GSS).
- access other product documentation.

Red Hat also hosts a large number of electronic mailing lists for discussion of Red Hat software and technology. You can find a list of publicly available mailing lists at https://www.redhat.com/mailman/listinfo. Click on the name of any mailing list to subscribe to that list or to access the list archives.

2.2. We Need Feedback!

If you find a typographical error in this manual, or if you have thought of a way to make this manual better, we would love to hear from you! Please submit a report in Bugzilla: http://bugzilla.redhat.com/ against the product **Red Hat Enterprise Virtualization Manager**.

When submitting a bug report, be sure to mention the manual's identifier: *Guides-Install*

If you have a suggestion for improving the documentation, try to be as specific as possible when describing it. If you have found an error, please include the section number and some of the surrounding text so we can find it easily.

Part I. Before you Begin

Chapter 1. Introduction

1.1. Red Hat Enterprise Virtualization Architecture

A Red Hat Enterprise Virtualization environment consists of:

- Virtual machine **hosts** using the Kernel-based Virtual Machine (KVM).
- **Agents and tools** running on hosts including VDSM, QEMU, and libvirt. These tools provide local management for virtual machines, networks and storage.
- The **Red Hat Enterprise Virtualization Manager**; a centralized management platform for the Red Hat Enterprise Virtualization environment. It provides a graphical interface where you can view, provision and manage resources.
- **Storage domains** to hold virtual resources like virtual machines, templates, ISOs.
- A **database** to track the state of and changes to the environment.
- Access to an external **Directory Server** to provide users and authentication.
- **Networking** to link the environment together. This includes physical network links, and logical networks.

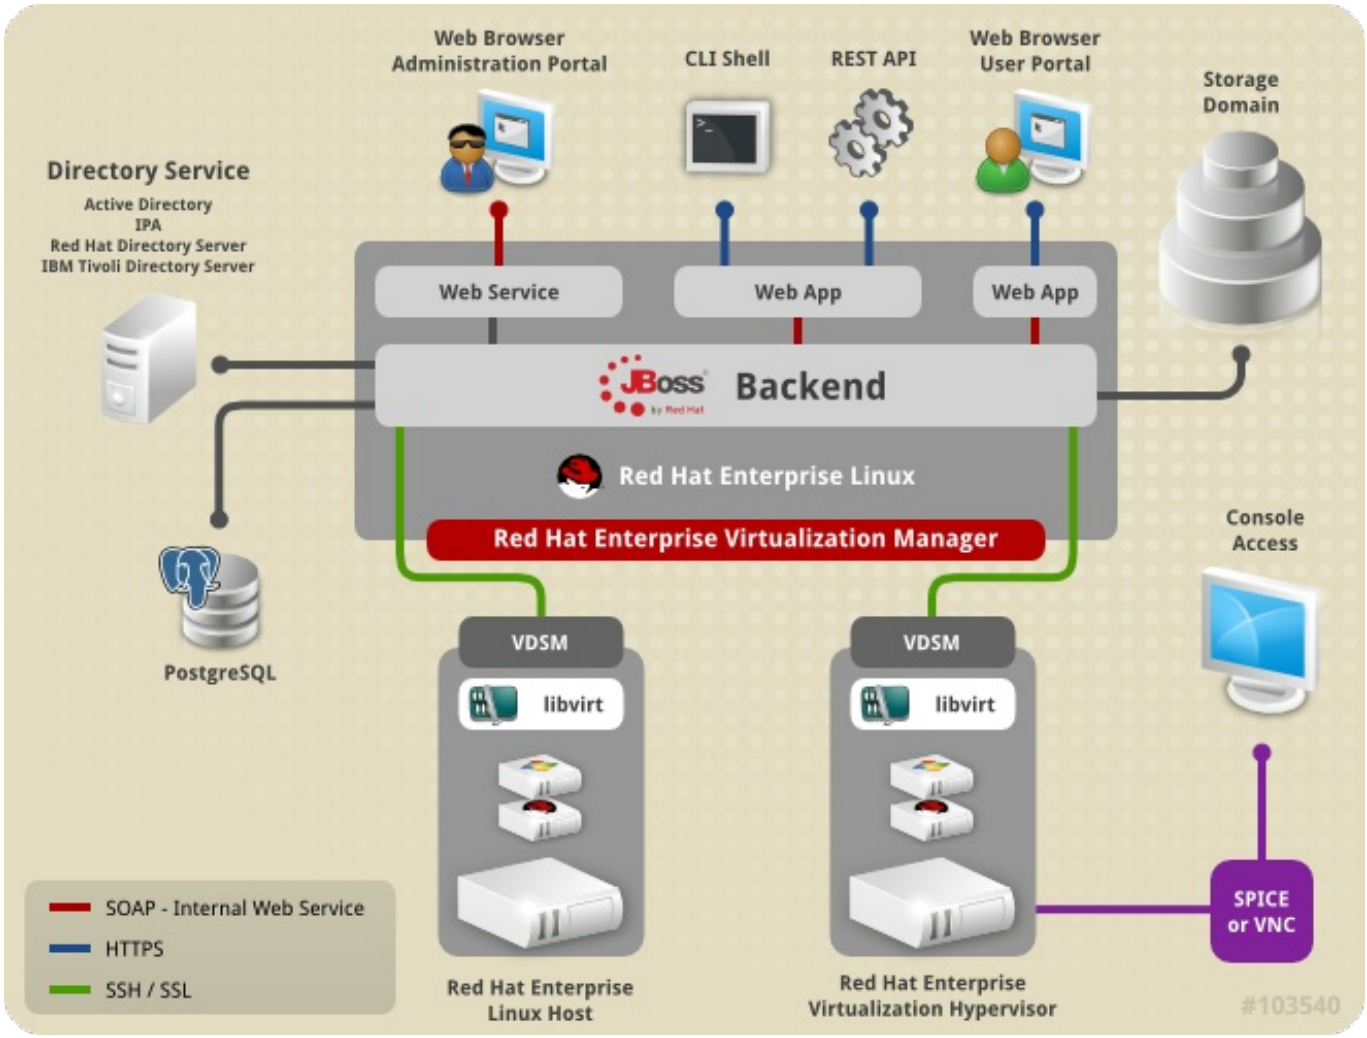

Figure 1.1. Red Hat Enterprise Virtualization Platform Overview

Report a bug

1.2. Red Hat Enterprise Virtualization System Components

The Red Hat Enterprise Virtualization version 3.3 environment consists of one or more hosts (either Red Hat Enterprise Linux 6.5 or later hosts or Red Hat Enterprise Virtualization Hypervisor 6.5 or later hosts) and at least one Red Hat Enterprise Virtualization Manager.

Hosts run virtual machines using KVM (Kernel-based Virtual Machine) virtualization technology.

The Red Hat Enterprise Virtualization Manager runs on a Red Hat Enterprise Linux 6 server and provides interfaces for controlling the Red Hat Enterprise Virtualization environment. It manages virtual machine and storage provisioning, connection protocols, user sessions, virtual machine images, and high availability virtual machines.

The Red Hat Enterprise Virtualization Manager is accessed through the Administration Portal using a web browser.

Report a bug

1.3. Red Hat Enterprise Virtualization Resources

The components of the Red Hat Enterprise Virtualization environment fall into two categories: physical resources, and logical resources. Physical resources are physical objects, such as host and storage servers. Logical resources are nonphysical groupings and processes, such as logical networks and virtual machine templates.

- **Data Center** - A data center is the highest level container for all physical and logical resources within a managed virtual environment. It is a collection of clusters, virtual machines, storage, and networks.
- **Clusters** - A cluster is a set of physical hosts that are treated as a resource pool for virtual machines. Hosts in a cluster share the same network infrastructure and storage. They form a migration domain within which virtual machines can be moved from host to host.
- **Logical Networks** - A logical network is a logical representation of a physical network. Logical networks group network traffic and communication between the Manager, hosts, storage, and virtual machines.
- **Hosts** - A host is a physical server that runs one or more virtual machines. Hosts are grouped into clusters. Virtual machines can be migrated from one host to another within a cluster.
- **Storage Pool** - The storage pool is a logical entity that contains a standalone image repository of a certain type, either iSCSI, Fibre Channel, NFS, or POSIX. Each storage pool can contain several domains, for storing virtual machine disk images, ISO images, and for the import and export of virtual machine images.
- **Virtual Machines** - A virtual machine is a virtual desktop or virtual server containing an operating system and a set of applications. Multiple identical virtual machines can be created in a **Pool**. Virtual machines are created, managed, or deleted by power users and accessed by users.
- **Template** - A template is a model virtual machine with predefined settings. A virtual machine that is based on a particular template acquires the settings of the template. Using templates is the quickest way of creating a large number of virtual machines in a single step.

- **Virtual Machine Pool** - A virtual machine pool is a group of identical virtual machines that are available on demand by each group member. Virtual machine pools can be set up for different purposes. For example, one pool can be for the Marketing department, another for Research and Development, and so on.
- **Snapshot** - A snapshot is a view of a virtual machine's operating system and all its applications at a point in time. It can be used to save the settings of a virtual machine before an upgrade or installing new applications. In case of problems, a snapshot can be used to restore the virtual machine to its original state.
- **User Types** - Red Hat Enterprise Virtualization supports multiple levels of administrators and users with distinct levels of permissions. System administrators can manage objects of the physical infrastructure, such as data centers, hosts, and storage. Users access virtual machines available from a virtual machine pool or standalone virtual machines made accessible by an administrator.
- **Events and Monitors** - Alerts, warnings, and other notices about activities help the administrator to monitor the performance and status of resources.
- **Reports** - A range of reports either from the reports module based on JasperReports, or from the data warehouse. Preconfigured or ad hoc reports can be generated from the reports module. Users can also generate reports using any query tool that supports SQL from a data warehouse that collects monitoring data for hosts, virtual machines, and storage.

Report a bug

1.4. Red Hat Enterprise Virtualization API Support Statement

Red Hat Enterprise Virtualization exposes a number of interfaces for interacting with the components of the virtualization environment. These interfaces are in addition to the user interfaces provided by the Red Hat Enterprise Virtualization Manager Administration, User, and Reports Portals. Many of these interfaces are fully supported. Some however are supported only for read access or only when your use of them has been explicitly requested by Red Hat Support.

Supported Interfaces for Read and Write Access

Direct interaction with these interfaces is supported and encouraged for both read and write access:

Representational State Transfer (REST) API

The REST API exposed by the Red Hat Enterprise Virtualization Manager is a fully supported interface for interacting with Red Hat Enterprise Virtualization Manager.

Software Development Kit (SDK)

The SDK provided by the *rhevm-sdk* package is a fully supported interface for interacting with Red Hat Enterprise Virtualization Manager.

Command Line Shell

The command line shell provided by the *rhevm-cli* package is a fully supported interface for interacting with the Red Hat Enterprise Virtualization Manager.

VDSM Hooks

The creation and use of VDSM hooks to trigger modification of virtual machines based on custom properties specified in the Administration Portal is supported on Red Hat Enterprise Linux virtualization hosts. The use of VDSM Hooks on virtualization hosts running Red Hat Enterprise Virtualization Hypervisor is not currently supported.

Supported Interfaces for Read Access

Direct interaction with these interfaces is supported and encouraged only for read access. Use of these interfaces for write access is not supported unless explicitly requested by Red Hat Support:

Red Hat Enterprise Virtualization Manager History Database

Read access to the Red Hat Enterprise Virtualization Manager history database using the database views specified in the Administration Guide is supported. Write access is *not* supported.

Libvirt on Virtualization Hosts

Read access to **`libvirt`** using the **`virsh -r`** command is a supported method of interacting with virtualization hosts. Write access is *not* supported.

Unsupported Interfaces

Direct interaction with these interfaces is *not* supported unless your use of them is explicitly requested by Red Hat Support:

The vdsClient Command

Use of the **`vdsClient`** command to interact with virtualization hosts is *not* supported unless explicitly requested by Red Hat Support.

Red Hat Enterprise Virtualization Hypervisor Console

Console access to Red Hat Enterprise Virtualization Hypervisor outside of the provided text user interface for configuration is *not* supported unless explicitly requested by Red Hat Support.

Red Hat Enterprise Virtualization Manager Database

Direct access to and manipulation of the Red Hat Enterprise Virtualization Manager database is *not* supported unless explicitly requested by Red Hat Support.

Important

Red Hat Support will not debug user created scripts or hooks except where it can be demonstrated that there is an issue with the interface being used rather than the user created script itself. For more general information about Red Hat support policies see https://access.redhat.com/support/offerings/production/soc.html.

Report a bug

1.5. Introduction to Virtual Machines

A virtual machine is a software implementation of a computer. The Red Hat Enterprise Virtualization environment enables you to create virtual desktops and virtual servers.

Virtual machines consolidate computing tasks and workloads. In traditional computing environments, workloads usually run on individually administered and upgraded servers. Virtual machines reduce the amount of hardware and administration required to run the same computing tasks and workloads.

Report a bug

1.6. Supported Virtual Machine Operating Systems

The operating systems that can be virtualized as guest operating systems in Red Hat Enterprise Virtualization are as follows:

Table 1.1. Operating systems that can be used as guest operating systems

Operating System	Architecture	SPICE support
Red Hat Enterprise Linux 3	32-bit, 64-bit	Yes
Red Hat Enterprise Linux 4	32-bit, 64-bit	Yes
Red Hat Enterprise Linux 5	32-bit, 64-bit	Yes
Red Hat Enterprise Linux 6	32-bit, 64-bit	Yes
SUSE Linux Enterprise Server 10 (select **Other Linux** for the guest type in the user interface)	32-bit, 64-bit	No
SUSE Linux Enterprise Server 11 (SPICE drivers (QXL) are not supplied by Red Hat. However, the distribution's vendor may provide spice drivers as part of their distribution.)	32-bit, 64-bit	No
Ubuntu 12.04 (Precise Pangolin LTS)	32-bit, 64-bit	Yes
Ubuntu 12.10 (Quantal Quetzal)	32-bit, 64-bit	Yes
Ubuntu 13.04 (Raring Ringtail)	32-bit, 64-bit	No
Ubuntu 13.10 (Saucy Salamander)	32-bit, 64-bit	Yes
Windows XP Service Pack 3 and newer	32-bit	Yes
Windows 7	32-bit, 64-bit	Yes
Windows 8	32-bit, 64-bit	No
Windows Server 2003 Service Pack 2 and newer	32-bit, 64-bit	Yes
Windows Server 2003 R2	32-bit, 64-bit	Yes
Windows Server 2008	32-bit, 64-bit	Yes
Windows Server 2008 R2	64-bit	Yes
Windows Server 2012	64-bit	No
Windows Server 2012 R2	64-bit	No

Of the operating systems that can be virtualized as guest operating systems in Red Hat Enterprise Virtualization, the operating systems that are supported by Global Support Services are as follows:

Table 1.2. Guest operating systems that are supported by Global Support Services

Operating System	Architecture
Red Hat Enterprise Linux 3	32-bit, 64-bit
Red Hat Enterprise Linux 4	32-bit, 64-bit
Red Hat Enterprise Linux 5	32-bit, 64-bit
Red Hat Enterprise Linux 6	32-bit, 64-bit
SUSE Linux Enterprise Server 10 (select **Other Linux** for the guest type in the user interface)	32-bit, 64-bit

Operating System	Architecture
SUSE Linux Enterprise Server 11 (SPICE drivers (QXL) are not supplied by Red Hat. However, the distribution's vendor may provide spice drivers as part of their distribution.)	32-bit, 64-bit
Windows XP Service Pack 3 and newer	32-bit
Windows 7	32-bit, 64-bit
Windows 8	32-bit, 64-bit
Windows Server 2003 Service Pack 2 and newer	32-bit, 64-bit
Windows Server 2003 R2	32-bit, 64-bit
Windows Server 2008	32-bit, 64-bit
Windows Server 2008 R2	64-bit
Windows Server 2012	64-bit
Windows Server 2012 R2	64-bit

Remote Desktop Protocol (RDP) is the default connection protocol for accessing Windows 8 and Windows 2012 guests from the user portal as Microsoft introduced changes to the Windows Display Driver Model that prevent SPICE from performing optimally.

Note

While Red Hat Enterprise Linux 3 and Red Hat Enterprise Linux 4 are supported, virtual machines running the 32-bit version of these operating systems cannot be shut down gracefully from the administration portal because there is no ACPI support in the 32-bit x86 kernel. To terminate virtual machines running the 32-bit version of Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4, right-click the virtual machine and select the **Power Off** option.

Note

See http://www.redhat.com/resourcelibrary/articles/enterprise-linux-virtualization-support for information about up-to-date guest support.

Report a bug

1.7. Red Hat Enterprise Virtualization Installation Workflow

Red Hat Enterprise Virtualization requires installation and configuration of several components to create a functioning virtualization environment. You must install and configure each component in the order shown in this checklist:

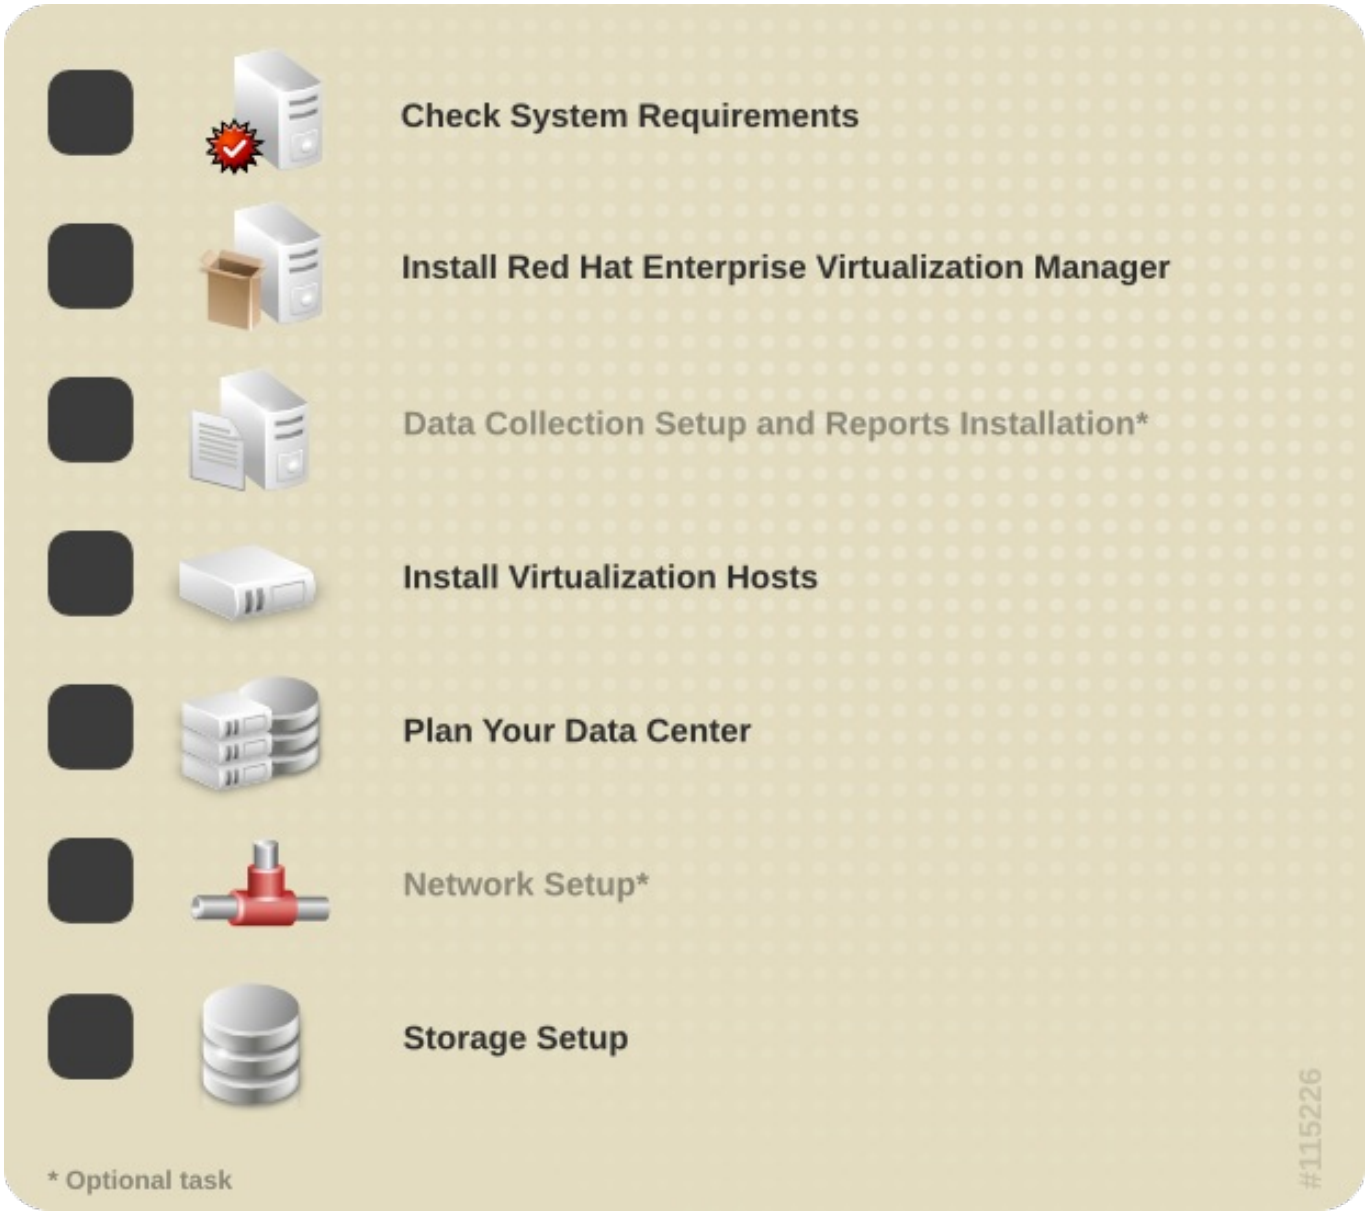

By completing these steps you will create a functioning Red Hat Enterprise Virtualization environment ready to host and run virtual machines.

- Chapter 2, *System Requirements*
- Part II, “Installing Red Hat Enterprise Virtualization Manager”
- Chapter 5, *Data Collection Setup and Reports Installation*
- Part III, “Installing Virtualization Hosts”
- Chapter 10, *Planning your Data Center*
- Chapter 11, *Network Setup*
- Chapter 12, *Storage Setup*

Report a bug

Chapter 2. System Requirements

2.1. Workflow Progress — System Requirements

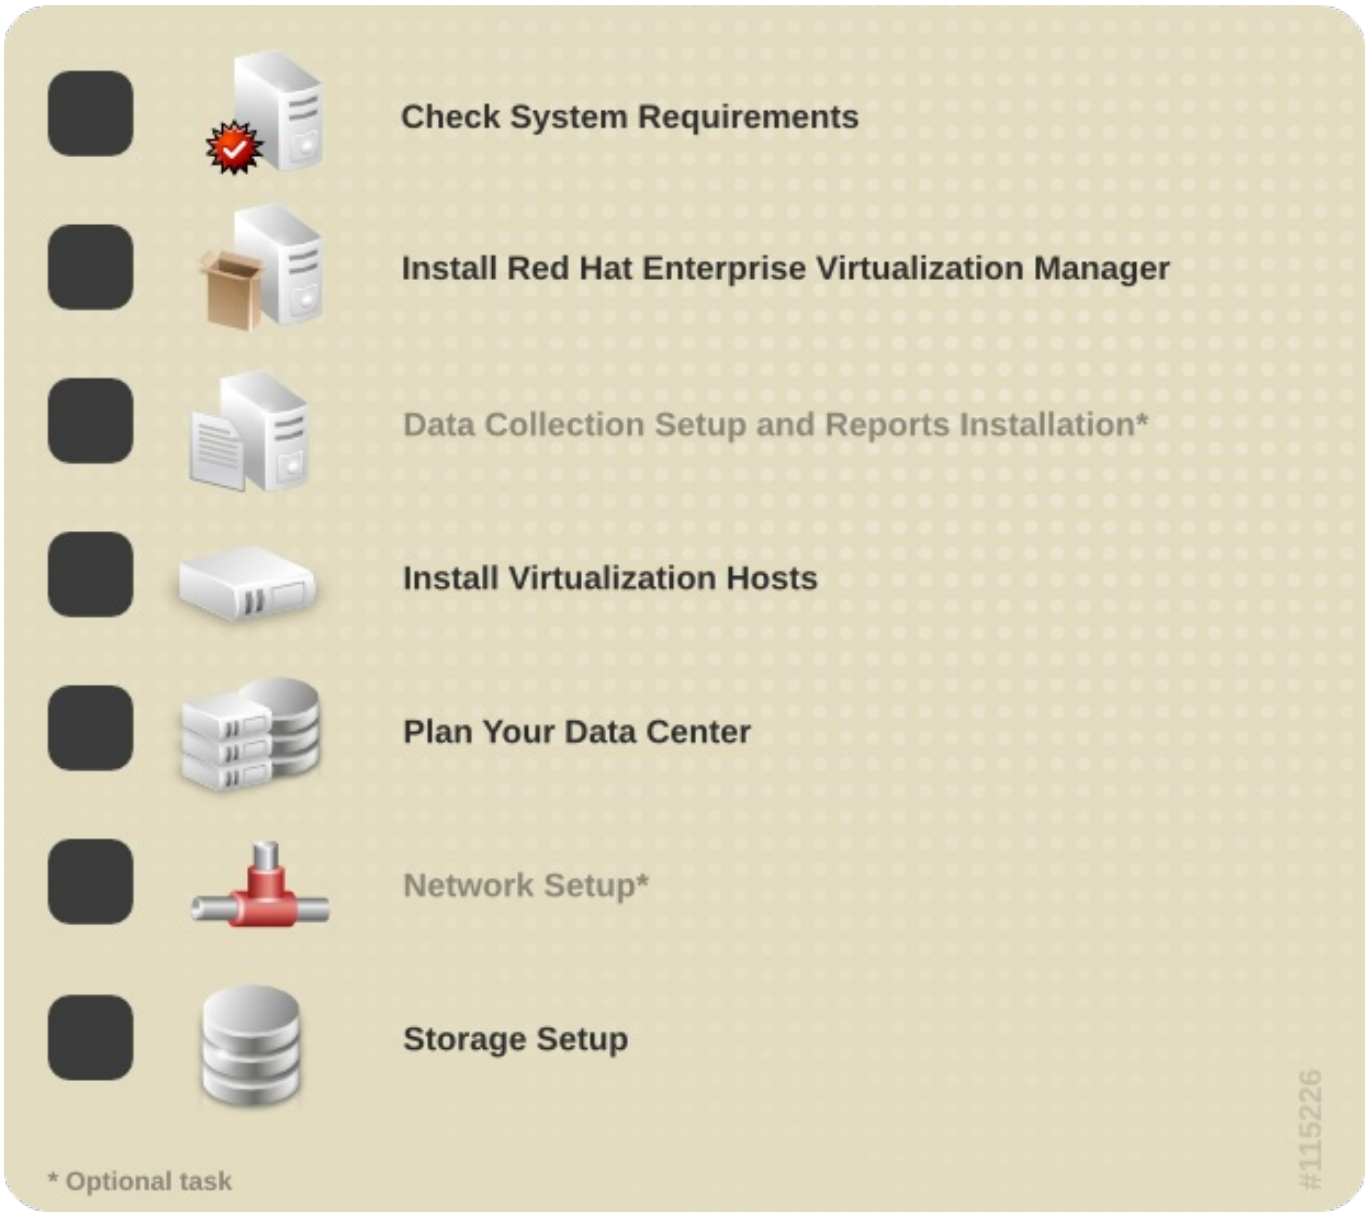

Report a bug

2.2. Hardware Requirements

2.2.1. Red Hat Enterprise Virtualization Hardware Requirements Overview

This section outlines the minimum hardware required to install, configure, and operate a Red Hat Enterprise Virtualization environment. To setup a Red Hat Enterprise Virtualization environment it is necessary to have, at least:

- one machine to act as the management server,
- one or more machines to act as virtualization hosts - at least two are required to support migration and power management,
- one or more machines to use as clients for accessing the Administration Portal.

- storage infrastructure provided by NFS, POSIX, iSCSI, SAN, or local storage.

The hardware required for each of these systems is further outlined in the following sections. The Red Hat Enterprise Virtualization environment also requires storage infrastructure that is accessible to the virtualization hosts. Storage infrastructure must be accessible using NFS, iSCSI, FC, or locally attached to virtualization hosts. The use of other POSIX compliant filesystems is also supported.

See Also:

- Section 2.2.2, “Red Hat Enterprise Virtualization Manager Hardware Requirements”
- Section 2.2.3, “Virtualization Host Hardware Requirements”

Report a bug

2.2.2. Red Hat Enterprise Virtualization Manager Hardware Requirements

The minimum and recommended hardware requirements outlined here are based on a typical small to medium sized installation. The exact requirements vary between deployments based on sizing and load. Please use these recommendations as a guide only.

Minimum

- A dual core CPU.
- 4 GB of available system RAM that is not being consumed by existing processes.
- 25 GB of locally accessible, writeable, disk space.
- 1 Network Interface Card (NIC) with bandwidth of at least 1 Gbps.

Recommended

- A quad core CPU or multiple dual core CPUs.
- 16 GB of system RAM.
- 50 GB of locally accessible, writeable, disk space.
- 1 Network Interface Card (NIC) with bandwidth of at least 1 Gbps.

The Red Hat Enterprise Virtualization Manager runs on Red Hat Enterprise Linux. To confirm whether or not specific hardware items are certified for use with Red Hat Enterprise Linux refer to http://www.redhat.com/rhel/compatibility/hardware/.

Report a bug

2.2.3. Virtualization Host Hardware Requirements

2.2.3.1. Virtualization Host Hardware Requirements Overview

Red Hat Enterprise Virtualization Hypervisors and Red Hat Enterprise Linux Hosts have a number of hardware requirements and supported limits.

- Section 2.2.3.2, “Virtualization Host CPU Requirements”
- Section 2.2.3.3, “Virtualization Host RAM Requirements”

- Section 2.2.3.4, “Virtualization Host Storage Requirements”
- Section 2.2.3.5, “Virtualization Host PCI Device Requirements”

Report a bug

2.2.3.2. Virtualization Host CPU Requirements

Red Hat Enterprise Virtualization supports the use of these CPU models in virtualization hosts:

- AMD Opteron G1
- AMD Opteron G2
- AMD Opteron G3
- AMD Opteron G4
- AMD Opteron G5
- Intel Conroe
- Intel Penryn
- Intel Nehalem
- Intel Westmere
- Intel Sandybridge
- Intel Haswell

All CPUs must have support for the Intel® 64 or AMD64 CPU extensions, and the AMD-V™ or Intel VT® hardware virtualization extensions enabled. Support for the **No eXecute** flag (NX) is also required. To check that your processor supports the required flags, and that they are enabled:

- At the Red Hat Enterprise Linux or Red Hat Enterprise Virtualization Hypervisor boot screen press any key and select the **Boot** or **Boot with serial console** entry from the list. Press **Tab** to edit the kernel parameters for the selected option. After the last kernel parameter listed ensure there is a **Space** and append the **rescue** parameter.
- Press **Enter** to boot into rescue mode.
- At the prompt which appears, determine that your processor has the required extensions and that they are enabled by running this command:

  ```
  # grep -E 'svm|vmx' /proc/cpuinfo | grep nx
  ```

  If any output is shown, the processor is hardware virtualization capable. If no output is shown it is still possible that your processor supports hardware virtualization. In some circumstances manufacturers disable the virtualization extensions in the BIOS. Where you believe this to be the case consult the system's BIOS and the motherboard manual provided by the manufacturer.

- As an additional check, verify that the **kvm** modules are loaded in the kernel:

  ```
  # lsmod | grep kvm
  ```

  If the output includes **kvm_intel** or **kvm_amd** then the **kvm** hardware virtualization modules are loaded and your system meets requirements.

Report a bug

2.2.3.3. Virtualization Host RAM Requirements

It is recommended that virtualization hosts have at least 2 GB of RAM. The amount of RAM required varies depending on:

- guest operating system requirements,
- guest application requirements, and
- memory activity and usage of guests.

The fact that KVM is able to over-commit physical RAM for virtualized guests must also be taken into account. This allows provisioning of guests with RAM requirements greater than physically present on the basis where not all guests will be at peak load concurrently. KVM does this by only allocating RAM for guests as required and shifting underutilized guests into swap.

A maximum of 2 TB of RAM per virtualization host is currently supported.

Report a bug

2.2.3.4. Virtualization Host Storage Requirements

Virtualization hosts require local storage to store configuration, logs, kernel dumps, and for use as swap space. The minimum storage requirements of the Red Hat Enterprise Virtualization Hypervisor are documented in this section. The storage requirements for Red Hat Enterprise Linux hosts vary based on the amount of disk space used by their existing configuration but are expected to be greater than those of the Red Hat Enterprise Virtualization Hypervisor.

It is recommended that each virtualization host has at least 2 GB of internal storage. The minimum supported internal storage for each Hypervisor is the total of that required to provision the following partitions:

- The root partitions require at least 512 MB of storage.
- The configuration partition requires at least 8 MB of storage.
- The recommended minimum size of the logging partition is 2048 MB.
- The data partition requires at least 256 MB of storage. Use of a smaller data partition may prevent future upgrades of the Hypervisor from the Red Hat Enterprise Virtualization Manager. By default all disk space remaining after allocation of swap space will be allocated to the data partition.
- The swap partition requires at least 8 MB of storage. The recommended size of the swap partition varies depending on both the system the Hypervisor is being installed upon and the anticipated level of overcommit for the environment. Overcommit allows the Red Hat Enterprise Virtualization environment to present more RAM to guests than is actually physically present. The default overcommit ratio is **0.5**.

  The recommended size of the swap partition can be determined by:

  - Multiplying the amount of system RAM by the expected overcommit ratio, and adding
  - 2 GB of swap space for systems with 4 GB of RAM or less, or
  - 4 GB of swap space for systems with between 4 GB and 16 GB of RAM, or
  - 8 GB of swap space for systems with between 16 GB and 64 GB of RAM, or

- 16 GB of swap space for systems with between 64 GB and 256 GB of RAM.

> **Example 2.1. Calculating Swap Partition Size**
>
> For a system with 8 GB of RAM this means the formula for determining the amount of swap space to allocate is:
>
> ```
> (8 GB x 0.5) + 4 GB = 8 GB
> ```

Important

By default the Red Hat Enterprise Virtualization Hypervisor defines a swap partition sized using the recommended formula. An overcommit ratio of **0.5** is used for this calculation. For some systems the result of this calculation may be a swap partition that requires more free disk space than is available at installation. Where this is the case Hypervisor installation will fail.

If you encounter this issue, manually set the sizes for the Hypervisor disk partitions using the ***storage_vol*** boot parameter.

> **Example 2.2. Manually Setting Swap Partition Size**
>
> In this example the ***storage_vol*** boot parameter is used to set a swap partition size of 4096 MB. Note that no sizes are specified for the other partitions, allowing the Hypervisor to use the default sizes.
>
> ```
> storage_vol=:4096::::
> ```

Important

The Red Hat Enterprise Virtualization Hypervisor does not support installation on **fakeraid** devices. Where a **fakeraid** device is present it must be reconfigured such that it no longer runs in RAID mode.

1. Access the RAID controller's BIOS and remove all logical drives from it.
2. Change controller mode to be non-RAID. This may be referred to as compatibility or JBOD mode.

Access the manufacturer provided documentation for further information related to the specific device in use.

Report a bug

2.2.3.5. Virtualization Host PCI Device Requirements

Virtualization hosts must have at least one network interface with a minimum bandwidth of 1 Gbps. It is recommended that each virtualization host have two network interfaces with a minimum bandwidth of 1 Gbps to support network intensive activity, including virtual machine migration.

Report a bug

2.3. Software Requirements

2.3.1. Red Hat Enterprise Virtualization Operating System Requirements

- Red Hat Enterprise Virtualization Manager requires Red Hat Enterprise Linux 6.5 Server. Complete successful installation of the operating system prior to commencing installation of the Red Hat Enterprise Virtualization Manager.

Important

See the Red Hat Enterprise Linux 6 Security Guide for security hardening information for your Red Hat Enterprise Linux Servers.

- Virtualization hosts must run either:
  - Red Hat Enterprise Virtualization Hypervisor 6.5
  - Red Hat Enterprise Linux 6.5

Important

Red Hat Enterprise Virtualization Manager must be installed on a base installation of Red Hat Enterprise Linux. Do not install any additional packages after the base installation, as they may cause dependency issues when attempting to install the packages required by the Manager.

Report a bug

2.3.2. Red Hat Enterprise Virtualization Manager Client Requirements

Use a client with a supported web browser to access the Administration Portal, and the User Portal. The portals support the following clients and browsers:

- Mozilla Firefox 17, and later, on Red Hat Enterprise Linux is required to access both portals.
- Internet Explorer 8, and later, on Microsoft Windows is required to access the User Portal. Use the desktop version, not the touchscreen version of Internet Explorer 10.
- Internet Explorer 9, and later, on Microsoft Windows is required to access the Administration Portal. Use the desktop version, not the touchscreen version of Internet Explorer 10.

Install a supported SPICE client to access virtual machine consoles. Supported SPICE clients are available on the following operating systems:

- Red Hat Enterprise Linux 5.8+ (i386, AMD64 and Intel 64)
- Red Hat Enterprise Linux 6.2+ (i386, AMD64 and Intel 64)
- Red Hat Enterprise Linux 6.5+ (i386, AMD64 and Intel 64)

- Windows XP
- Windows XP Embedded (XPe)
- Windows 7 (x86, AMD64 and Intel 64)
- Windows 8 (x86, AMD64 and Intel 64)
- Windows Embedded Standard 7
- Windows 2008/R2 (x86, AMD64 and Intel 64)
- Windows Embedded Standard 2009
- Red Hat Enterprise Virtualization Certified Linux-based thin clients

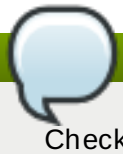

Note

Check the Red Hat Enterprise Virtualization Manager Release Notes to see which SPICE features your client supports.

When you access the portal(s) using Mozilla Firefox the SPICE client is provided by the *spice-xpi* package, which you must manually install using **yum**.

When you access the portal(s) using Internet Explorer the SPICE ActiveX control will automatically be downloaded and installed.

Report a bug

2.3.3. Red Hat Enterprise Virtualization Manager Software Channels

Installation of the Red Hat Enterprise Virtualization Manager requires that the system be subscribed to a number of Red Hat Network channels in addition to those required for Red Hat Enterprise Linux. These channels are used to retrieve both the initial installation packages and later updates as they become available.

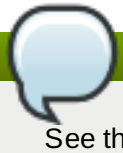

Note

See the *Red Hat Enterprise Virtualization Manager Release Notes* for specific channel names current to your system.

You must ensure that you have entitlements to the required channels listed here before proceeding with installation.

Certificate-based Red Hat Network

- The **Red Hat Enterprise Linux Server** entitlement, provides Red Hat Enterprise Linux.
- The **Red Hat Enterprise Virtualization** entitlement, provides Red Hat Enterprise Virtualization Manager.
- The **Red Hat JBoss Enterprise Application Platform** entitlement, provides the supported release of the application platform on which the Manager runs.

Red Hat Network Classic

- The **Red Hat Enterprise Linux Server (v. 6 for 64-bit x86_64)** channel, also referred to as **rhel-x86_64-server-6**, provides Red Hat Enterprise Linux 6 Server. The Channel Entitlement name for this channel is **Red Hat Enterprise Linux Server (v. 6)**.
- The **RHEL Server Supplementary (v. 6 64-bit x86_64)** channel, also referred to as **rhel-x86_64-server-supplementary-6**, provides the *virtio-win* package. The *virtio-win* package provides the Windows VirtIO drivers for use in virtual machines. The Channel Entitlement Name for the supplementary channel is **Red Hat Enterprise Linux Server Supplementary (v. 6)**.
- The **Red Hat Enterprise Virtualization Manager (v3.3 x86_64)** channel, also referred to as **rhel-x86_64-server-6-rhevm-3.3**, provides Red Hat Enterprise Virtualization Manager. The Channel Entitlement Name for this channel is **Red Hat Enterprise Virtualization Manager (v3)**.
- The **Red Hat JBoss EAP (v 6) for 6Server x86_64** channel, also referred to as **jbappplatform-6-x86_64-server-6-rpm**, provides the supported release of the application platform on which the Manager runs. The Channel Entitlement Name for this channel is **Red Hat JBoss Enterprise Application Platform (v 4, zip format)**.

See Also:

- Section 3.3, “Subscribing to the Red Hat Enterprise Virtualization Channels”

Report a bug

2.3.4. Directory Services

2.3.4.1. About Directory Services

The term directory service refers to the collection of software, hardware, and processes that store information about an enterprise, subscribers, or both, and make that information available to users. A directory service consists of at least one directory server and at least one directory client program. Client programs can access names, phone numbers, addresses, and other data stored in the directory service.

See Also:

- Section 2.3.4.2, “Directory Services Support in Red Hat Enterprise Virtualization”

Report a bug

2.3.4.2. Directory Services Support in Red Hat Enterprise Virtualization

During installation Red Hat Enterprise Virtualization Manager creates its own internal administration user, **admin**. This account is intended for use when initially configuring the environment, and for troubleshooting. To add other users to Red Hat Enterprise Virtualization you will need to attach a directory server to the Manager using the Domain Management Tool, **engine-manage-domains**.

Once at least one directory server has been attached to the Manager you will be able to add users that exist in the directory server and assign roles to them using the Administration Portal. Users will be identified by their User Principal Name (UPN) of the form ***user@domain***. Attachment of more than one directory server to the Manager is also supported.

The directory servers supported for use with Red Hat Enterprise Virtualization 3.3 are:

- Active Directory
- Identity Management (IdM)
- Red Hat Directory Server 9 (RHDS 9)
- OpenLDAP

You must ensure that the correct DNS records exist for your directory server. In particular you must ensure that the DNS records for the directory server include:

- A valid pointer record (PTR) for the directory server's reverse look-up address.
- A valid service record (SRV) for LDAP over TCP port **389**.
- A valid service record (SRV) for Kerberos over TCP port **88**.
- A valid service record (SRV) for Kerberos over UDP port **88**.

If these records do not exist in DNS then you will be unable to add the domain to the Red Hat Enterprise Virtualization Manager configuration using **`engine-manage-domains`**.

For more detailed information on installing and configuring a supported directory server, see the vendor's documentation:

- Active Directory - http://technet.microsoft.com/en-us/windowsserver/dd448614.
- Identity Management (IdM) - http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Identity_Management_Guide/index.html
- Red Hat Directory Server (RHDS) - http://docs.redhat.com/docs/en-US/Red_Hat_Directory_Server/index.html
- OpenLDAP - http://www.openldap.org/doc/

Important

A user must be created in the directory server specifically for use as the Red Hat Enterprise Virtualization administrative user. Do *not* use the administrative user for the directory server as the Red Hat Enterprise Virtualization administrative user.

Important

It is not possible to install Red Hat Enterprise Virtualization Manager (*rhevm*) and IdM (*ipa-server*) on the same system. IdM is incompatible with the *mod_ssl* package, which is required by Red Hat Enterprise Virtualization Manager.

Important

If you are using Active Directory as your directory server, and you wish to use **sysprep** in the creation of Templates and Virtual Machines, then the Red Hat Enterprise Virtualization administrative user must be delegated control over the Domain to:

- **Join a computer to the domain**
- **Modify the membership of a group**

For information on creation of user accounts in Active Directory, see http://technet.microsoft.com/en-us/library/cc732336.aspx.

For information on delegation of control in Active Directory, see http://technet.microsoft.com/en-us/library/cc732524.aspx.

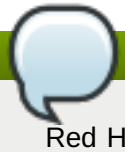

Note

Red Hat Enterprise Virtualization Manager uses Kerberos to authenticate with directory servers. RHDS does not provide native support for Kerberos. If you are using RHDS as your directory server then you must ensure that the directory server is made a service within a valid Kerberos domain. To do this you will need to perform these steps while referring to the relevant directory server documentation:

- Configure the **memberOf** plug-in for RHDS to allow group membership. In particular ensure that the value of the ***memberofgroupattr*** attribute of the **memberOf** plug-in is set to **uniqueMember**. In **OpenLDAP**, the **memberOf** functionality is not called a "plugin". It is called an "overlay" and requires no configuration after installation.

  Consult the Red Hat Directory Server 9.0 *Plug-in Guide* for more information on configuring the **memberOf** plug-in.

- Define the directory server as a service of the form **ldap/*hostname*@*REALMNAME*** in the Kerberos realm. Replace *hostname* with the fully qualified domain name associated with the directory server and *REALMNAME* with the fully qualified Kerberos realm name. The Kerberos realm name must be specified in capital letters.
- Generate a **keytab** file for the directory server in the Kerberos realm. The **keytab** file contains pairs of Kerberos principals and their associated encrypted keys. These keys will allow the directory server to authenticate itself with the Kerberos realm.

  Consult the documentation for your Kerberos principle for more information on generating a **keytab** file.

- Install the **keytab** file on the directory server. Then configure RHDS to recognize the **keytab** file and accept Kerberos authentication using GSSAPI.

  Consult the Red Hat Directory Server 9.0 *Administration Guide* for more information on configuring RHDS to use an external **keytab** file.

- Test the configuration on the directory server by using the **kinit** command to authenticate as a user defined in the Kerberos realm. Once authenticated run the **ldapsearch** command against the directory server. Use the ***-Y GSSAPI*** parameters to ensure the use of Kerberos for authentication.

See Also:

- Section B.1, “Domain Management Tool”
- Section 2.3.5.3, “Directory Server Firewall Requirements”

Report a bug

2.3.5. Firewall Configuration

2.3.5.1. Red Hat Enterprise Virtualization Manager Firewall Requirements

The Red Hat Enterprise Virtualization Manager requires that a number of ports be opened to allow network traffic through the system's firewall. The **engine-setup** script is able to configure the firewall automatically, but this will overwrite any pre-existing firewall configuration.

Where an existing firewall configuration exists the firewall rules required by the Manager must instead be manually inserted into it. The **engine-setup** command will save a list of the **iptables** rules required in the **/usr/share/ovirt-engine/conf/iptables.example** file.

The firewall configuration documented here assumes a default configuration. Where non-default HTTP and HTTPS ports are chosen during installation adjust the firewall rules to allow network traffic on the ports that were selected - not the default ports (**80** and **443**) listed here.

Table 2.1. Red Hat Enterprise Virtualization Manager Firewall Requirements

Port(s)	Protocol	Source	Destination	Purpose
-	ICMP	- Red Hat Enterprise Virtualization Hypervisor(s) - Red Hat Enterprise Linux host(s)	- Red Hat Enterprise Virtualization Manager	When registering to the Red Hat Enterprise Virtualization Manager, virtualization hosts send an ICMP ping request to the Manager to confirm that it is online.
22	TCP	- System(s) used for maintenance of the Manager including backend configuration, and software upgrades.	- Red Hat Enterprise Virtualization Manager	SSH (optional)

Port(s)	Protocol	Source	Destination	Purpose
80, 443	TCP	• Administration Portal clients • User Portal clients • Red Hat Enterprise Virtualization Hypervisor(s) • Red Hat Enterprise Linux host(s) • REST API clients	• Red Hat Enterprise Virtualization Manager	Provides HTTP and HTTPS access to the Manager.

Important

In environments where the Red Hat Enterprise Virtualization Manager is also required to export NFS storage, such as an ISO Storage Domain, additional ports must be allowed through the firewall. Grant firewall exceptions for the ports applicable to the version of NFS in use:

NFSv4

- TCP port **2049** for NFS.

NFSv3

- TCP and UDP port **2049** for NFS.
- TCP and UDP port **111** (**rpcbind**/**sunrpc**).
- TCP and UDP port specified with **MOUNTD_PORT="*port*"**
- TCP and UDP port specified with **STATD_PORT="*port*"**
- TCP port specified with **LOCKD_TCPPORT="*port*"**
- UDP port specified with **LOCKD_UDPPORT="*port*"**

The **MOUNTD_PORT**, **STATD_PORT**, **LOCKD_TCPPORT**, and **LOCKD_UDPPORT** ports are configured in the **/etc/sysconfig/nfs** file.

Report a bug

2.3.5.2. Virtualization Host Firewall Requirements

Both Red Hat Enterprise Linux hosts and Red Hat Enterprise Virtualization Hypervisors require that a number of ports be opened to allow network traffic through the system's firewall. In the case of the Red Hat Enterprise Virtualization Hypervisor these firewall rules are configured automatically. For Red Hat Enterprise Linux hosts however it is necessary to manually configure the firewall.

Table 2.2. Virtualization Host Firewall Requirements

Port(s)	Protocol	Source	Destination	Purpose
22	TCP	• Red Hat Enterprise Virtualization Manager	• Red Hat Enterprise Virtualization Hypervisor(s) • Red Hat Enterprise Linux host(s)	Secure Shell (SSH) access.
5900 - 6411	TCP	• Administration Portal clients • User Portal clients	• Red Hat Enterprise Virtualization Hypervisor(s) • Red Hat Enterprise Linux host(s)	Remote guest console access via VNC and SPICE. These ports must be open to facilitate client access to virtual machines.
5989	TCP, UDP	• Common Information Model Object Manager (CIMOM)	• Red Hat Enterprise Virtualization Hypervisor(s) • Red Hat Enterprise Linux host(s)	Used by Common Information Model Object Managers (CIMOM) to monitor virtual machines running on the virtualization host. If you wish to use a CIMOM to monitor the virtual machines in your virtualization environment then you must ensure that this port is open.
16514	TCP	• Red Hat Enterprise Virtualization Hypervisor(s) • Red Hat Enterprise Linux host(s)	• Red Hat Enterprise Virtualization Hypervisor(s) • Red Hat Enterprise Linux host(s)	Virtual machine migration using **`libvirt`**.

Port(s)	Protocol	Source	Destination	Purpose
49152 - 49216	TCP	• Red Hat Enterprise Linux Hypervisor(s) • Red Hat Enterprise Linux host(s)	• Red Hat Enterprise Linux Hypervisor(s) • Red Hat Enterprise Linux host(s)	Virtual machine migration and fencing using VDSM. These ports must be open facilitate both automated and manually initiated migration of virtual machines.
54321	TCP	• Red Hat Enterprise Virtualization Manager • Red Hat Enterprise Virtualization Hypervisor(s) • Red Hat Enterprise Linux host(s)	• Red Hat Enterprise Virtualization Hypervisor(s) • Red Hat Enterprise Linux host(s)	VDSM communications with the Manager and other virtualization hosts.

Example 2.3. Option Name: IPTablesConfig

Recommended (default) values: Automatically generated by vdsm bootstrap script

```
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [0:0]
-A INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT
-A INPUT -p icmp -j ACCEPT
-A INPUT -i lo -j ACCEPT
# vdsm
-A INPUT -p tcp --dport 54321 -j ACCEPT
# libvirt tls
-A INPUT -p tcp --dport 16514 -j ACCEPT
# SSH
-A INPUT -p tcp --dport 22 -j ACCEPT
# guest consoles
-A INPUT -p tcp -m multiport --dports 5634:6166 -j ACCEPT
# migration
-A INPUT -p tcp -m multiport --dports 49152:49216 -j ACCEPT
# snmp
-A INPUT -p udp --dport 161 -j ACCEPT
# Reject any other input traffic
```

```
-A INPUT -j REJECT --reject-with icmp-host-prohibited
-A FORWARD -m physdev ! --physdev-is-bridged -j REJECT --reject-with
icmp-host-prohibited
COMMIT
```

Report a bug

2.3.5.3. Directory Server Firewall Requirements

Red Hat Enterprise Virtualization requires a directory server to support user authentication. A number of ports must be opened in the directory server's firewall to support GSS-API authentication as used by the Red Hat Enterprise Virtualization Manager.

Table 2.3. Host Firewall Requirements

Port(s)	Protocol	Source	Destination	Purpose
88, 464	TCP, UDP	Red Hat Enterprise Virtualization Manager	Directory server	Kerberos authentication.
389, 636	TCP	Red Hat Enterprise Virtualization Manager	Directory server	Lightweight Directory Access Protocol (LDAP) and LDAP over SSL.

Report a bug

2.3.5.4. Database Server Firewall Requirements

Red Hat Enterprise Virtualization supports the use of a remote database server. If you plan to use a remote database server with Red Hat Enterprise Virtualization then you must ensure that the remote database server allows connections from the Manager.

Table 2.4. Host Firewall Requirements

Port(s)	Protocol	Source	Destination	Purpose
5432	TCP, UDP	Red Hat Enterprise Virtualization Manager	PostgreSQL database server	Default port for PostgreSQL database connections.

If you plan to use a local database server on the Manager itself, which is the default option provided during installation, then no additional firewall rules are required.

Report a bug

2.3.6. Required User Accounts and Groups

2.3.6.1. Red Hat Enterprise Virtualization Manager User Accounts

When the *rhevm* package is installed a number of user accounts are created to support Red Hat Enterprise Virtualization. The user accounts created as a result of *rhevm* package installation are as follows. The default user identifier (UID) for each account is also provided:

- The **vdsm** user (UID **36**). Required for support tools that mount and access NFS storage domains.
- The **ovirt** user (UID **108**). Owner of the **ovirt-engine** Red Hat JBoss Enterprise Application Platform instance.

Report a bug

2.3.6.2. Red Hat Enterprise Virtualization Manager Groups

When the *rhevm* package is installed a number of user groups are created. The user groups created as a result of *rhevm* package installation are as follows. The default group identifier (GID) for each group is also listed:

- The **kvm** group (GID **36**). Group members include:
  - The **vdsm** user.
- The **ovirt** group (GID **108**). Group members include:
  - The **ovirt** user.

Report a bug

2.3.6.3. Virtualization Host User Accounts

When the *vdsm* and *qemu-kvm-rhev* packages are installed a number of user accounts are created. These are the user accounts that are created on the virtualization host as a result of *vdsm* and *qemu-kvm-rhev* package installation. The default user identifier (UID) for each entry is also listed:

- The **vdsm** user (UID **36**).
- The **qemu** user (UID **107**).
- The **sanlock** user (UID **179**).

In addition Red Hat Enterprise Virtualization Hypervisor hosts define a **admin** user (UID **500**). This **admin** user is not created on Red Hat Enterprise Linux virtualization hosts. The **admin** user is created with the required permissions to run commands as the **root** user using the **sudo** command. The **vdsm** user which is present on both types of virtualization hosts is also given access to the **sudo** command.

Important

The user identifiers (UIDs) and group identifiers (GIDs) allocated may vary between systems. The **vdsm** user however is fixed to a UID of **36** and the **kvm** group is fixed to a GID of **36**.

If UID **36** or GID **36** is already used by another account on the system then a conflict will arise during installation of the *vdsm* and *qemu-kvm-rhev* packages.

Report a bug

2.3.6.4. Virtualization Host Groups

When the *vdsm* and *qemu-kvm-rhev* packages are installed a number of user groups are created. These are the groups that are created on the virtualization host as a result of *vdsm* and *qemu-kvm-rhev* package installation. The default group identifier (GID) for each entry is also listed:

- The **kvm** group (GID **36**). Group members include:
  - The **qemu** user.
  - The **sanlock** user.
- The **qemu** group (GID **107**). Group members include:
  - The **vdsm** user.
  - The **sanlock** user.

Important

The user identifiers (UIDs) and group identifiers (GIDs) allocated may vary between systems. The **vdsm** user however is fixed to a UID of **36** and the **kvm** group is fixed to a GID of **36**.

If UID **36** or GID **36** is already used by another account on the system then a conflict will arise during installation of the *vdsm* and *qemu-kvm-rhev* packages.

Report a bug

Part II. Installing Red Hat Enterprise Virtualization Manager

Chapter 3. Manager Installation

3.1. Workflow Progress — Installing Red Hat Enterprise Virtualization Manager

Report a bug

3.2. Installing the Red Hat Enterprise Virtualization Manager

Prerequisites:

- Section 2.2, “Hardware Requirements”
- Section 2.3, “Software Requirements”

Overview

The Red Hat Enterprise Virtualization Manager can be installed under one of two arrangements - a standard setup in which the Manager is installed on an independent physical machine or virtual machine, or a self-hosted engine setup in which the Manager runs on a virtual machine that the Manager itself controls.

Important

While the prerequisites for and basic configuration of the Red Hat Enterprise Virtualization Manager itself are the same for both standard and self-hosted engine setups, the process for setting up a self-hosted engine is different from that of a standard setup. To install the Manager as a self-hosted engine, follow the instructions in *Self-Hosted Engine*.

Prerequisites

Before installing the Red Hat Virtualization Manager, you must ensure you meet all the prerequisites. To complete installation of the Red Hat Enterprise Virtualization Manager successfully, you must also be able to determine:

1. The ports to be used for HTTP and HTTPS communication. The defaults ports are **80** and **443** respectively.
2. The fully qualified domain name (FQDN) of the system on which the Manager is to be installed.
3. The password you will use to secure the Red Hat Enterprise Virtualization administration account.
4. The location of the database server to be used. You can use the setup script to install and configure a local database server or use an existing remote database server. To use a remote database server you will need to know:
   - The host name of the system on which the remote database server exists.
   - The port on which the remote database server is listening.
   - That the **uuid-ossp** extension had been loaded by the remote database server.

   You must also know the user name and password of a user that is known to the remote database server. The user must have permission to create databases in PostgreSQL.
5. The organization name to use when creating the Manager's security certificates.
6. The storage type to be used for the initial data center attached to the Manager. The default is NFS.
7. The path to use for the ISO share, if the Manager is being configured to provide one. The display name, which will be used to label the domain in the Red Hat Enterprise Virtualization Manager also needs to be provided.
8. The firewall rules, if any, present on the system that need to be integrated with the rules required for the Manager to function.

Configuration

Before installation is completed the values selected are displayed for confirmation. Once the values have been confirmed they are applied and the Red Hat Enterprise Virtualization Manager is ready for use.

Example 3.1. Completed Installation

```
--== CONFIGURATION PREVIEW ==--

Database name                       : engine
Database secured connection         : False
Database host                       : localhost
Database user name                  : engine
Database host name validation       : False
Database port                       : 5432
NFS setup                           : True
PKI organization                    : demo.redhat.com
Application mode                    : both
Firewall manager                    : iptables
Update Firewall                     : True
Configure WebSocket Proxy           : True
Host FQDN                           : rhevm33.demo.redhat.com
NFS mount point                     : /var/lib/exports/iso
Datacenter storage type             : nfs
Configure local database            : True
Set application as default page     : True
Configure Apache SSL                : True

Please confirm installation settings (OK, Cancel) [OK]:
```

Note

Automated installations are created by providing **engine-setup** with an answer file. An answer file contains answers to the questions asked by the setup command.

- To create an answer file, use the ***--generate-answer*** parameter to specify a path and file name with which to create the answer file. When this option is specified, the **engine-setup** command records your answers to the questions in the setup process to the answer file.

  ```
  # engine-setup --generate-answer=ANSWER_FILE
  ```

- To use an answer file for a new installation, use the ***--config-append*** parameter to specify the path and file name of the answer file to be used. The **engine-setup** command will use the answers stored in the file to complete the installation.

  ```
  # engine-setup --config-append=ANSWER_FILE
  ```

Run **engine-setup *--help*** for a full list of parameters.

Note

Offline installation requires the creation of a software repository local to your Red Hat Enterprise Virtualization environment. This software repository must contain all of the packages required to install Red Hat Enterprise Virtualization Manager, Red Hat Enterprise Linux virtualization hosts, and Red Hat Enterprise Linux virtual machines. To create such a repository, see the *Installing Red Hat Enterprise Virtualization Offline* technical brief, available at https://access.redhat.com/knowledge/techbriefs/installing-red-hat-enterprise-virtualization-offline-rhev-30.

Report a bug

3.3. Subscribing to the Red Hat Enterprise Virtualization Channels

3.3.1. Subscribing to the Red Hat Enterprise Virtualization Manager Channels using Subscription Manager

Summary

To install Red Hat Enterprise Virtualization Manager, you must first register the target system in Red Hat Network and subscribe to the required software channels.

Procedure 3.1. Subscribing to the Red Hat Enterprise Virtualization Manager Channels using Subscription Manager

1. **Register the System with Subscription Manager**

   Run the **subscription-manager register** command to register the system with Red Hat Network. To complete registration successfully, you will need to supply your Red Hat Network **Username** and **Password** when prompted.

   ```
   # subscription-manager register
   ```

2. **Identify Available Entitlement Pools**

   To subscribe the system to Red Hat Enterprise Virtualization, you must locate the identifiers for the relevant entitlement pools. Use the ***list*** action of the **subscription-manager** to find these.

   To identify available subscription pools for **Red Hat Enterprise Linux Server**, use the command:

   ```
   # subscription-manager list --available | grep -A8 "Red Hat
   Enterprise Linux Server"
   ```

   To identify available subscription pools for **Red Hat Enterprise Virtualization**, use the command:

   ```
   # subscription-manager list --available | grep -A8 "Red Hat
   Enterprise Virtualization"
   ```

3. **Attach Entitlement Pools to the System**

   Using the pool identifiers located in the previous step, attach the **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** entitlements to the system. To do so, use the ***attach*** parameter of the **subscription-manager** command, replacing *[POOLID]* with each of the pool identifiers:

   ```
   # subscription-manager attach --pool=[POOLID]
   ```

4. **Enable the Red Hat Enterprise Virtualization Manager 3.3 Repository**

   Attaching a **Red Hat Enterprise Virtualization** entitlement pool also subscribes the system to the Red Hat Enterprise Virtualization Manager 3.3 software repository. By default, this software repository is available but disabled. The Red Hat Enterprise Virtualization Manager 3.3 software repository must be enabled using the **yum-config-manager** command:

   ```
   # yum-config-manager --enable rhel-6-server-rhevm-3.3-rpms
   ```

5. **Enable the Supplementary Repository**

   Attaching a **Red Hat Enterprise Linux Server** entitlement pool also subscribes the system to the supplementary software repository. By default, this software repository is available but disabled. The supplementary software repository must be enabled using the **yum-config-manager** command:

   ```
   # yum-config-manager --enable rhel-6-server-supplementary-rpms
   ```

6. **Enable the Red Hat JBoss Enterprise Application Platform Repository**

   The JBoss Enterprise Application Platform channels required for Red Hat Enterprise Virtualization are included in the Red Hat Enterprise Virtualization subscription. However, the repository that contains these channels is disabled by default, and must be enabled using the **yum-config-manager** command:

   ```
   # yum-config-manager --enable jb-eap-6-for-rhel-6-server-rpms
   ```

Result

The system is now registered with Red Hat Network and subscribed to the channels required for Red Hat Enterprise Virtualization Manager installation.

Report a bug

3.3.2. Subscribing to the Red Hat Enterprise Virtualization Manager Channels Using RHN Classic

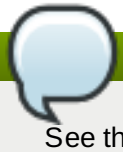

Note

See the *Red Hat Enterprise Virtualization Manager Release Notes* for specific channel names current to your system.

Summary

To install Red Hat Enterprise Virtualization Manager you must first register the target system to Red Hat Network and subscribe to the required software channels.

Procedure 3.2. Subscribing to the Red Hat Enterprise Virtualization Manager Channels using RHN Classic

1. Run the **rhn_register** command to register the system with Red Hat Network. To complete registration successfully you will need to supply your Red Hat Network user name and password. Follow the onscreen prompts to complete registration of the system.

   ```
   # rhn_register
   ```

2. **Subscribe to Required Channels**

   You must subscribe the system to the required channels using either the web interface to Red Hat Network or the command line **rhn-channel** command.

   A. **Using the rhn-channel Command**

   Run the **rhn-channel** command to subscribe the system to each of the required channels. The commands which need to be run are:

   ```
   # rhn-channel --add --channel=rhel-x86_64-server-6
   # rhn-channel --add --channel=rhel-x86_64-server-supplementary-6
   # rhn-channel --add --channel=rhel-x86_64-server-6-rhevm-3.3
   # rhn-channel --add --channel=jbappplatform-6-x86_64-server-6-
   rpm
   ```

Important

If you are not the administrator for the machine as defined in Red Hat Network, or the machine is not registered to Red Hat Network, then use of the **rhn-channel** command will result in an error:

```
Error communicating with server. The message was:
Error Class Code: 37
Error Class Info: You are not allowed to perform
administrative tasks on this system.
Explanation:
     An error has occurred while processing your request. If
this problem
     persists please enter a bug report at
bugzilla.redhat.com.
     If you choose to submit the bug report, please be sure
to include
     details of what you were trying to do when this error
occurred and
     details on how to reproduce this problem.
```

If you encounter this error when using **rhn-channel** then to add the channel to the system you must use the web user interface.

B. **Using the Web Interface to Red Hat Network**

To add a channel subscription to a system from the web interface:

a. Log on to Red Hat Network (http://rhn.redhat.com).

b. Move the mouse cursor over the **Subscriptions** link at the top of the screen, and then click the **Registered Systems** link in the menu that appears.

c. Select the system to which you are adding channels from the list presented on the screen, by clicking the name of the system.

d. Click **Alter Channel Subscriptions** in the **Subscribed Channels** section of the screen.

e. Select the channels to be added from the list presented on the screen. Red Hat Enterprise Virtualization Manager requires:

- The **Red Hat Enterprise Linux Server (v. 6 for 64-bit x86_64)** channel. This channel is located under the **Release Channels for Red Hat Enterprise Linux 6 for x86_64** expandable menu.
- The **RHEL Server Supplementary (v. 6 64-bit x86_64)** channel. This channel is located under the **Release Channels for Red Hat Enterprise Linux 6 for x86_64** expandable menu.
- The **Red Hat Enterprise Virtualization Manager (v.3.3 x86_64)** channel. This channel is located under the **Additional Services Channels for Red Hat Enterprise Linux 6 for x86_64** expandable menu.

- The **Red Hat JBoss EAP (v 6) for 6Server x86_64** channel. This channel is located under the **Additional Services Channels for Red Hat Enterprise Linux 6 for x86_64** expandable menu.

f. Click the **Change Subscription** button to finalize the change.

Result

The system is now registered with Red Hat Network and subscribed to the channels required for Red Hat Enterprise Virtualization Manager installation.

Report a bug

3.4. Installing the Red Hat Enterprise Virtualization Manager Packages

Prerequisites:

- Section 3.3, “Subscribing to the Red Hat Enterprise Virtualization Channels”

Summary

The Red Hat Enterprise Virtualization Manager and all of its dependencies are available via Red Hat Network. To install the Manager you must ensure connectivity with either Red Hat Network or a Red Hat Network Satellite Server is available. You must also ensure that you have access to all required Red Hat Network channels.

Procedure 3.3. Installing the Red Hat Enterprise Virtualization Manager Packages

1. Use **yum** to ensure that the most up to date versions of all installed packages are in use.

```
# yum upgrade
```

2. Use **yum** to initiate installation of the *rhevm* package and all dependencies. You must run this command as the **root** user.

```
# yum install rhevm
```

Note

Installing the *rhevm* package also installs all packages which it depends on. This includes the *java-1.7.0-openjdk* package. The *java-1.7.0-openjdk* package provides the OpenJDK Java Virtual Machine (JVM) required to run Red Hat Enterprise Virtualization Manager.

3. The *rhevm* package includes the *rhevm-doc* package as a dependency. The *rhevm-doc* package provides a local copy of the Red Hat Enterprise Virtualization documentation suite. This documentation is also used to provide context sensitive help links from the Administration and User Portals.

   As localized versions of this package become available they will be released to Red Hat Network. Follow these steps to find and install any available localized Red Hat Enterprise Virtualization documentation packages that you require:

a. Use the **yum** command to search for translated Red Hat Enterprise Virtualization Manager documentation packages:

```
# yum search rhevm-doc
```

b. While logged in as the **root** user use the **yum** command to install translated Red Hat Enterprise Virtualization Manager documentation packages. Here the Japanese (**ja-JP**) version of the package is installed:

```
# yum install rhevm-doc-ja-JP
```

Result

All required packages and dependencies are installed. The required Java runtime environment was also selected as the default for the system. You must now configure the Red Hat Enterprise Virtualization Manager.

See Also:

- Section 3.5, “Configuring Red Hat Enterprise Virtualization Manager”
- Section 6.1.2, “Updating Red Hat Enterprise Virtualization Manager”
- Section 3.10, “Removing Red Hat Enterprise Virtualization Manager”

Report a bug

3.5. Configuring Red Hat Enterprise Virtualization Manager

Prerequisites:

- Section 3.4, “Installing the Red Hat Enterprise Virtualization Manager Packages”

When installation of the packages is complete, the Red Hat Enterprise Virtualization Manager must be configured. The **engine-setup** script is provided to assist with this task. This script asks you a series of questions, and configures your environment based on your answers. After the required values have been provided, the updated configuration is applied and the Red Hat Enterprise Virtualization Manager services are started.

The **engine-setup** script guides you through several distinct configuration stages, each comprising several steps that require user input. At each step, suggested configuration defaults are provided in square brackets. When these default values are acceptable for a given step, you can press the **Enter** key to accept the default values and proceed to the next step or stage.

Procedure 3.4. Manager Configuration Overview

1. **Packages Check**

   The **engine-setup** script checks to see if it is performing an upgrade or an installation, and whether any updates are available for the packages linked to the Manager. No user input is required at this stage.

```
[ INFO  ] Checking for product updates...
[ INFO  ] No product updates found
```

2. **Network Configuration**

   A reverse lookup is performed on your host name, which is automatically detected. You can correct the auto-detected host name if it is incorrect, or if you are using virtual hosts. Your fully-qualified domain name should have both forward and reverse lookup records in DNS, especially if will also install the reports server.

   ```
   Host fully qualified DNS name of this server [autodetected host
   name]:
   ```

   The **engine-setup** script checks your firewall configuration, and offers to modify it for you to open the ports used by the Manager for external communications (for example, TCP ports 80 and 443). If you do not allow **engine-setup** script to modify your iptables configuration, you must manually open the ports used by the Red Hat Enterprise Virtualization Manager (see *Red Hat Enterprise Virtualization Manager Firewall Requirements*).

   ```
   iptables was detected on your computer. Do you wish Setup to
   configure it? (yes, no) [yes]:
   ```

3. **Database Configuration**

   You can use either a local or remote **Postgres** database. The **engine-setup** script can configure your database completely automatically, including adding a user and a database, or use values that you supply.

   ```
   Where is the database located? (Local, Remote) [Local]:
   Setup can configure the local postgresql server automatically for
   the engine to run. This may conflict with existing applications.
   Would you like Setup to automatically configure postgresql, or
   prefer to perform that manually? (Automatic, Manual) [Automatic]:
   ```

4. **OVirt Engine Configuration**

   Set a password for the automatically created administrative user of the Red Hat Enterprise Virtualization Manager: **admin@internal**.

   ```
   Engine admin password:
   Confirm engine admin password:
   ```

   Select **Gluster**, **Virtualization**, or **Both**. **Both** gives the greatest flexibility.

   ```
   Data center (Both, Virt, Gluster) [Both]:
   ```

   Choose the initial data center storage type. You can have many data centers in your environment, each with a different type of storage. Here, you are choosing the storage type of your first data center.

   ```
   Default storage type: (NFS, FC, ISCSI, POSIXFS) [NFS]:
   ```

5. **PKI Configuration**

   The Manager uses certificates to communicate securely with its hosts. You provide the organization name for the certificate. This certificate can also optionally be used to secure https communications with the Manager.

```
Organization name for certificate [autodetected domain-based name]:
```

6. **Apache Configuration**

   The Red Hat Enterprise Virtualization Manager uses the Apache web server to present a landing page to users. The **engine-setup** script can make the landing page of the Manager the default page presented by Apache.

```
Setup can configure the default page of the web server to present
the application home page. This may conflict with existing
applications.
Do you wish to set the application as the default web page of the
server? (Yes, No) [Yes]:
```

   By default, external ssl (https) communications with the Manager are secured with the self-signed certificate created in the PKI configuration stage to securely communicate with hosts. Another certificate may be chosen for external https connections, without affecting how the Manager communicates with hosts.

```
Setup can configure apache to use SSL using a certificate issued
from the internal CA
Do you wish Setup to configure that, or prefer to perform that
manually? (Automatic, Manual) [Automatic]:
```

7. **System Configuration**

   The **engine-setup** script can create an NFS export on the Manager to use as an ISO storage domain. Hosting the ISO domain locally to the Manager simplifies keeping some elements of your environment up to date.

```
Configure an NFS share on this server to be used as an ISO Domain?
(Yes, No) [Yes]:
Local ISO domain path [/var/lib/exports/iso]:
Local ISO domain name [ISO_DOMAIN]:
```

8. **Websocket Proxy Server Configuration**

   The **engine-setup** script can optionally configure a websocket proxy server for allowing users to connect to virtual machines via the noVNC or HTML 5 consoles.

```
Configure WebSocket Proxy on this machine? (Yes, No) [Yes]:
```

9. **End of Configuration**

   The **engine-setup** script validates all of your answers, and warns you of any possible problem with them. User input is only required if some of the answers you provided may adversely impact your environment.

```
          --== END OF CONFIGURATION ==--
Would you like transactions from the Red Hat Access Plugin sent
from the RHEV Manager to be brokered through a proxy server? (Yes,
No) [No]:
```

```
[ INFO  ] Stage: Setup validation
```

10. **Preview, and Summary**

    During the preview phase, the **engine-setup** scripts shows you the configuration values you have entered, and gives you the opportunity to change your mind. If you choose to proceed, **engine-setup** configures your Red Hat Enterprise Virtualization Manager installation based on the answers you provided in the configuration stages.

```
          --== CONFIGURATION PREVIEW ==--
Database name                       : engine
Database secured connection         : False
Database host                       : localhost
Database user name                  : engine
Database host name validation       : False
Database port                       : 5432
NFS setup                           : True
PKI organization                    : Your Org
NFS mount point                     : /var/lib/exports/iso
Application mode                    : both
Firewall manager                    : iptables
Configure WebSocket Proxy           : True
Host FQDN                           : Your Manager's FQDN
Datacenter storage type             : nfs
Configure local database            : True
Set application as default page     : True
Configure Apache SSL                : True

Please confirm installation settings (OK, Cancel) [OK]:
```

    When your environment is configured, the **engine-setup** script provides some details about accessing your environment and it's security details.

```
A default ISO NFS share has been created on this host.
If IP based access restrictions are required, edit: entry
/var/lib/exports/iso in /etc/exports
SSH fingerprint: 87:af:b5:fe:7a:e5:1b:64:83:57:02:07:62:eb:8c:18
Internal CA SHA1
Fingerprint=7B:DF:2A:EE:18:C8:B1:CC:F7:6B:59:42:A3:96:BC:44:32:98:F
F:A6
Web access is enabled at:
      http://manager.fqdn:80/ovirt-engine
      https://manager.fqdn:443/ovirt-engine
Please use the user "admin" and password specified in order to
login into oVirt Engine
```

11. **Clean up and Termination**

    The **engine-setup** script cleans up unnecessary files created during the configuration process, and outputs the location of the log file for the Red Hat Enterprise Virtualization Manager configuration process.

```
[ INFO  ] Stage: Clean up
Log file is located at /var/log/ovirt-engine/setup/ovirt-engine-
```

```
setup-installation-date.log
[ INFO  ] Stage: Pre-termination
[ INFO  ] Stage: Termination
[ INFO  ] Execution of setup completed successfully
```

Result

When the **engine-setup** script completes successfully, the Red Hat Enterprise Virtualization Manager is configured and running on your server. You can log in as the **admin@internal** user to continue configuring the Manager, by adding clusters, hosts, and more.

You may want to immediately link your Red Hat Enterprise Virtualization Manager to a directory server, so you can add additional users. Red Hat Enterprise Virtualization supports directory services from RHDS, IdM, and Active Directory. Add a directory server to your environment using the **engine-manage-domains** command.

The **engine-setup** script also saves the answers you gave during configuration to a file, to help with disaster recovery.

See Also:

- Section 2.3.5.1, "Red Hat Enterprise Virtualization Manager Firewall Requirements"
- Chapter 5, *Data Collection Setup and Reports Installation*
- Part III, "Installing Virtualization Hosts"

Report a bug

3.6. Passwords in Red Hat Enterprise Virtualization Manager

Passwords are generated by default in Red Hat Enterprise Virtualization in the following way. **engine-setup** generates a temporary password for the database and a temporary password for the **admin@internal** account. To change the **admin@internal** password, run **engine-config** on the Red Hat Enterprise Virtualization Manager.

It is possible to specify a password manually by using the command **engine-setup --answer-file=*/[path_to_answer_file]***, with the temporary password specified in the answer file. Answer files are generated with the **engine-setup --generate-answer=file** command and option. The format of the answer file is as follows:

```
# action=setup
[environment:default]
OVESETUP_CORE/engineStop=none:None
OVESETUP_DIALOG/confirmSettings=bool:True
OVESETUP_DB/database=str:engine
OVESETUP_DB/fixDbViolations=none:None
OVESETUP_DB/secured=bool:False
OVESETUP_DB/host=str:localhost
OVESETUP_DB/user=str:engine
OVESETUP_DB/securedHostValidation=bool:False
OVESETUP_DB/password=str:0056jKkY
OVESETUP_DB/port=int:5432
OVESETUP_SYSTEM/nfsConfigEnabled=bool:True
...
OVESETUP_APACHE/configureSsl=bool:True
```

```
OSETUP_RPMDISTRO/requireRollback=none:None
OSETUP_RPMDISTRO/enableUpgrade=none:None
OVESETUP_AIO/configure=none:None
OVESETUP_AIO/storageDomainDir=none:None
```

Report a bug

3.7. Preparing a PostgreSQL Database for Use with Red Hat Enterprise Virtualization Manager

Summary

You can manually configure a database server to host the database used by the Red Hat Enterprise Virtualization Manager. The database can be hosted either locally on the machine on which the Red Hat Enterprise Virtualization Manager is installed, or remotely on another machine on which **posgresql** is installed.

Complete these steps on your database server before running the **engine-setup** command.

Procedure 3.5. Preparing a PostgreSQL Database for use with Red Hat Enterprise Virtualization Manager

1. Run the following commands to initialize the **postgresql** database, start the **postgresql** service and ensure this service starts on boot:

```
# service postgresql initdb
# service postgresql start
# chkconfig postgresql on
```

2. Create a user for the Red Hat Enterprise Virtualization Manager to use when it writes to and reads from the database, and a database in which to store data about your environment. This step is required for both local and remote databases. Use the **psql** terminal as the **postgres** user.

```
# su - postgres
$ psql
postgres=# create role [user name] with login encrypted password
'[password]';
postgres=# create database [database name] owner [user name]
template template0 encoding 'UTF8' lc_collate 'en_US.UTF-8'
lc_ctype 'en_US.UTF-8';
```

3. Run the following commands to connect to the new database and add the **plpgsql** language:

```
postgres=# \c [database name]
CREATE LANGUAGE plpgsql;
```

4. Ensure the database can be accessed remotely by enabling client authentication. Edit the **/var/lib/pgsql/data/pg_hba.conf** file, and add the following in accordance with the location of the database:

   - For local databases, add the two following lines immediately underneath the line starting with **Local** at the bottom of the file:

```
host    [database name]    [user name]    0.0.0.0/0  md5
host    [database name]    [user name]    ::0/0      md5
```

- For remote databases, add the following line immediately underneath the line starting with **Local** at the bottom of the file, replacing *X.X.X.X* with the IP address of the Manager:

```
host    [database name]    [user name]    X.X.X.X/32   md5
```

5. Allow TCP/IP connections to the database. This step is required for remote databases. Edit the **/var/lib/pgsql/data/postgresql.conf** file, and add the following line:

```
listen_addresses='*'
```

This example configures the **postgresql** service to listen for connections on all interfaces. You can specify an interface by giving its IP address.

6. Restart the **postgres** service. This step is required on both local and remote manually configured database servers.

```
service postgresql restart
```

Result

You have manually configured a PostgreSQL database to use with the Red Hat Enterprise Virtualization Manager.

Report a bug

3.8. Configuring the Manager to Use a Manually Configured Local or Remote PostgreSQL Database

Prerequisites:

- Section 3.7, “Preparing a PostgreSQL Database for Use with Red Hat Enterprise Virtualization Manager”

Summary

During the database configuration stage of configuring the Red Hat Enterprise Virtualization Manager using the **engine-setup** script, you can choose to use a manually configured database. You can select to use a locally or remotely installed **PostgreSQL** database.

Procedure 3.6. Configuring the Manager to use a Manually Configured Local or Remote PostgreSQL Database

1. During configuration of the Red Hat Enterprise Virtualization Manager using the **engine-setup** script, you are prompted to decide where your database is located:

```
Where is the database located? (Local, Remote) [Local]:
```

The steps involved in manually configuring the Red Hat Enterprise Virtualization Manger to use local or remotely hosted databases are the same, except that to use a remotely hosted database, you must provide the host name of the remote database server and the port on which it is listening.

2. When prompted, enter **Manual** to manually configure the database:

```
Would you like Setup to automatically configure postgresql, or
prefer to perform that manually? (Automatic, Manual) [Automatic]:
Manual
```

3. If you are using a remotely hosted database, supply the **engine-setup** script with the host name of your database server and the port on which it is listening:

```
Database host [localhost]:
Database port [5432]:
```

4. For both local and remotely hosted databases, you are then prompted to confirm whether your database uses a secured connection, and for the name of the database you configured, the user the Manager will use to access the database, and the password of that user.

```
Database secured connection (Yes, No) [No]:
Database name [engine]:
Database user [engine]:
Database password:
```

Note

Using a secured connection to your database requires you to have also manually configured secured database connections.

Result

You have configured the Red Hat Enterprise Virtualization Manager to use a manually configured database. The **engine-setup** script continues with the rest of your environment configuration.

Report a bug

3.9. Connecting to the Administration Portal

Summary

Access the Administration Portal using a web browser.

Procedure 3.7. Connecting to the Administration Portal

1. Open a supported web browser on your client system.
2. Navigate to **https://*your-manager-fqdn*/ovirt-engine**, replacing *your-manager-fqdn* with the fully qualified domain name that you provided during installation.

   The first time that you connect, you are prompted to trust the certificate being used to secure communications between your browser and the web server.

3. The login screen is displayed. Enter your **User Name** and **Password** in the fields provided. If you are logging in for the first time, use the user name **admin** in conjunction with the administrator password that you specified during installation.

4. Select the directory services domain to authenticate against from the **Domain** list provided. If you are logging in using the internal **admin** user name, then select the **internal** domain.

5. The Administration Portal is available in multiple languages. The default selection will be chosen based on the locale settings of your web browser. If you would like to view the Administration Portal in a language other than that selected by default, select your preferred language from the list.

6. Click **Login** to log in.

Result

You have now successfully logged into the Administration Portal.

Report a bug

3.10. Removing Red Hat Enterprise Virtualization Manager

Summary

Red Hat Enterprise Virtualization Manager provides the **engine-cleanup** script to allow quick and easy removal of the files associated with the Red Hat Enterprise Virtualization Manager environment.

Procedure 3.8. Removing Red Hat Enterprise Virtualization Manager

1. Run the **engine-cleanup** command on the system on which Red Hat Enterprise Virtualization Manager is installed.

```
# engine-cleanup
```

2. You are prompted to confirm removal of all Red Hat Enterprise Virtualization Manager components. These components include PKI keys, the locally hosted ISO domain file system layout, PKI configuration, the local NFS exports configuration, and the Engine database content.

```
Do you want to remove all components? (Yes, No) [Yes]:
```

Note

A backup of the Engine database and a compressed archive of the PKI keys and configuration are always automatically created. These are saved under **/var/lib/ovirt-engine/backups/**, and include the date and **engine-** and **engine-pki-** in their file names respectively.

3. You are given another opportunity to change your mind and cancel the removal of the Red Hat Enterprise Virtualization Manager. If you choose to proceed, the **ovirt-engine** service is stopped, and your environment's configuration is removed in accordance with the options you selected.

```
During execution engine service will be stopped (OK, Cancel) [OK]:
ovirt-engine is about to be removed, data will be lost (OK, Cancel)
[Cancel]:OK
```

Result

The configuration files of your environment have been removed according to your selections when you ran **engine-cleanup**.

```
          --== SUMMARY ==--

A backup of the database is available at /var/lib/ovirt-
engine/backups/engine-date-and-extra-characters.sql
Engine setup successfully cleaned up
A backup of PKI configuration and keys is available at /var/lib/ovirt-
engine/backups/engine-pki-date-and-extra-characters.tar.gz

          --== END OF SUMMARY ==--

[ INFO  ] Generating answer file '/var/lib/ovirt-
engine/setup/answers/20130827181911-cleanup.conf'
[ INFO  ] Stage: Clean up
          Log file is located at /var/log/ovirt-engine/setup/ovirt-
engine-remove-date.log
[ INFO  ] Stage: Pre-termination
[ INFO  ] Stage: Termination
[ INFO  ] Execution of cleanup completed successfully
```

You can now safely remove the Red Hat Enterprise Virtualization packages using the **yum** command.

```
# yum remove rhevm* vdsm-bootstrap
```

Report a bug

Chapter 4. Self-Hosted Engine

4.1. About the Self-Hosted Engine

A self-hosted engine is a virtualized environment in which the engine, or Manager, runs on a virtual machine on the hosts managed by that engine. The virtual machine is created as part of the host configuration, and the engine is installed and configured in parallel to that host configuration process, referred to in these procedures as the deployment.

The virtual machine running the engine is created to be highly available. This means that if the host running the virtual machine goes into maintenance mode, or fails unexpectedly, the virtual machine will be migrated automatically to another host in the environment.

The primary benefit of the self-hosted engine is that it requires less hardware to deploy an instance of Red Hat Enterprise Virtualization as the engine runs as a virtual machine, not on physical hardware. Additionally, the engine is configured to be highly available automatically, rather than requiring a separate cluster.

The self-hosted engine currently only runs on Red Hat Enterprise Linux 6.5 hosts. Red Hat Enterprise Virtualization Hypervisors and older versions of Red Hat Enterprise Linux are not recommended for use with a self-hosted engine.

Report a bug

4.2. Limitations of the Self-Hosted Engine

At present there are two main limitations of the self-hosted engine configuration:

- An NFS storage domain is required for the configuration.
- The host and hosted engine must use Red Hat Enterprise Linux 6.5 or above. Red Hat Enterprise Virtualization Hypervisors are not supported.

Report a bug

4.3. Installing the Self-Hosted Engine

Summary

Install a Red Hat Enterprise Virtualization environment that takes advantage of the self-hosted engine feature, in which the engine is installed on a virtual machine within the environment itself.

You must be subscribed to the appropriate Red Hat Network channels to install the packages.

All steps in this procedure are to be conducted as the **root** user.

Important

While the *ovirt-hosted-engine-setup* package is provided by the Red Hat Enterprise Virtualization Manager channel and can be installed using the standard channels for the Manager, the *vdsm* package is a dependency of the *ovirt-hosted-engine-setup* package and is provided by the Red Hat Enterprise Virt Management Agent channel, which must be enabled. This channel is **rhel-6-server-rhev-mgmt-agent-rpms** in Subscription Manager and **rhel-x86_64-rhev-mgmt-agent-6** in RHN Classic.

Procedure 4.1. Installing the Self-Hosted Engine

1. Use **yum** to ensure that the most up-to-date versions of all installed packages are in use.

```
# yum upgrade
```

2. Use **yum** to initiate installation of the *ovirt-hosted-engine-setup* package and all dependencies.

```
# yum install ovirt-hosted-engine-setup
```

Result

You have installed the *ovirt-hosted-engine-setup* package and are ready to configure the self-hosted engine.

Report a bug

4.4. Configuring the Self-Hosted Engine

Summary

When package installation is complete, the Red Hat Enterprise Virtualization Manager must be configured. The **hosted-engine** deployment script is provided to assist with this task. The script asks you a series of questions, and configures your environment based on your answers. When the required values have been provided, the updated configuration is applied and the Red Hat Enterprise Virtualization Manager services are started.

The **hosted-engine** deployment script guides you through several distinct configuration stages. The script suggests possible configuration defaults in square brackets. Where these default values are acceptable, no additional input is required.

This procedure requires a new Red Hat Enterprise Linux 6.5 host with the *ovirt-hosted-engine-setup* package installed. This host is referred to as 'Host-HE1', with a fully qualified domain name (FQDN) of **Host-HE1.example.com** in this procedure.

The hosted engine, the virtual machine created during configuration of Host-HE1 to manage the environment, is referred to as 'my-engine'. You will be prompted by the **hosted-engine** deployment script to access this virtual machine multiple times to install an operating system and to configure the engine.

All steps in this procedure are to be conducted as the **root** user for the specified machine.

Procedure 4.2. Configuring the Self-Hosted Engine

1. **Initiating Hosted Engine Deployment**

   Begin configuration of the self-hosted environment by deploying the **hosted-engine** customization script on Host_HE1. To escape the script at any time, use the **CTRL+D** keyboard combination to abort deployment.

   ```
   # hosted-engine --deploy
   ```

2. **Configuring Storage**

   Select the version of NFS and specify the full address, using either the FQDN or IP address, and path name of the shared storage domain. Choose the storage domain and storage data center names to be used in the environment.

   ```
   During customization use CTRL-D to abort.
   Please specify the storage you would like to use (nfs3, nfs4)[nfs3]:
   Please specify the full shared storage connection path to use
   (example: host:/path): storage.example.com:/hosted_engine/nfs
   [ INFO  ] Installing on first host
   Please provide storage domain name. [hosted_storage]:
   Local storage datacenter name is an internal name and currently
   will not be shown in engine's admin UI.Please enter local
   datacenter name [hosted_datacenter]:
   ```

3. **Configuring the Network**

   The script detects possible network interface controllers (NICs) to use as a management bridge for the environment. It then checks your firewall configuration and offers to modify it for console (SPICE or VNC) access HostedEngine-VM. Provide a pingable gateway IP address, to be used by the **ovirt-ha-agent** to help determine a host's suitability for running HostedEngine-VM.

   ```
   Please indicate a nic to set rhevm bridge on: (eth1, eth0) [eth1]:
   iptables was detected on your computer, do you wish setup to
   configure it? (Yes, No)[Yes]:
   Please indicate a pingable gateway IP address [X.X.X.X]:
   ```

4. **Configuring the Virtual Machine**

   The script creates a virtual machine to be configured as the Red Hat Enterprise Virtualization Manager, the hosted engine referred to in this procedure as HostedEngine-VM. Specify the boot device and, if applicable, the path name of the installation media, the CPU type, the number of virtual CPUs, and the disk size. Specify a MAC address for the HostedEngine-VM, or accept a randomly generated one. The MAC address can be used to update your DHCP server prior to installing the operating system on the virtual machine. Specify memory size and console connection type for the creation of HostedEngine-VM.

   ```
   Please specify the device to boot the VM from (cdrom, disk, pxe)
   [cdrom]:
   The following CPU types are supported by this host:
            - model_Penryn: Intel Penryn Family
            - model_Conroe: Intel Conroe Family
   Please specify the CPU type to be used by the VM [model_Penryn]:
   Please specify the number of virtual CPUs for the VM [Defaults to
   minimum requirement: 2]:
   ```

```
Please specify the disk size of the VM in GB [Defaults to minimum
requirement: 25]:
You may specify a MAC address for the VM or accept a randomly
generated default [00:16:3e:77:b2:a4]:
Please specify the memory size of the VM in MB [Defaults to minimum
requirement: 4096]:
Please specify the console type you would like to use to connect to
the VM (vnc, spice) [vnc]:
```

5. **Configuring the Hosted Engine**

   Specify the name for Host-HE1 to be identified in the Red Hat Enterprise Virtualization environment, and the password for the **admin@internal** user to access the Administrator Portal. Provide the FQDN for HostedEngine-VM; this procedure uses the FQDN *HostedEngine-VM.example.com*. Finally, provide the name and TCP port number of the SMTP server, the email address used to send email notifications, and a comma-separated list of email addresses to receive these notifications.

   ```
   Enter the name which will be used to identify this host inside the
   Administrator Portal [hosted_engine_1]: Host-HE1
   Enter 'admin@internal' user password that will be used for
   accessing the Administrator Portal:
   Confirm 'admin@internal' user password:
   Please provide the FQDN for the engine you would like to use. This
   needs to match the FQDN that you will use for the engine
   installation within the VM: HostedEngine-VM.example.com
   Please provide the name of the SMTP server through which we will
   send notifications [localhost]:
   Please provide the TCP port number of the SMTP server [25]:
   Please provide the email address from which notifications will be
   sent [root@localhost]:
   Please provide a comma-separated list of email addresses which will
   get notifications [root@localhost]:
   ```

6. **Configuration Preview**

   Before proceeding, the **hosted-engine** script displays the configuration values you have entered, and prompts for confirmation to proceed with these values.

   ```
   Bridge interface                   : eth1
   Engine FQDN                        : HostedEngine-VM.example.com
   Bridge name                        : rhevm
   SSH daemon port                    : 22
   Firewall manager                   : iptables
   Gateway address                    : X.X.X.X
   Host name for web application      : Host-HE1
   Host ID                            : 1
   Image size GB                      : 25
   Storage connection                 :
   storage.example.com:/hosted_engine/nfs
   Console type                       : vnc
   Memory size MB                     : 4096
   MAC address                        : 00:16:3e:77:b2:a4
   Boot type                          : pxe
   ```

```
Number of CPUs                     : 2
CPU Type                           : model_Penryn

Please confirm installation settings (Yes, No)[No]:
```

7. **Creating HostedEngine-VM**

   The script creates the virtual machine that will be configured to be HostedEngine-VM and provides connection details. You will need to install an operating system on HostedEngine-VM before the **`hosted-engine`** script can proceed on Host-HE1.

```
[ INFO  ] Generating answer file '/etc/ovirt-hosted-
engine/answers.conf'
[ INFO  ] Stage: Transaction setup
[ INFO  ] Stage: Misc configuration
[ INFO  ] Stage: Package installation
[ INFO  ] Stage: Misc configuration
[ INFO  ] Configuring libvirt
[ INFO  ] Generating VDSM certificates
[ INFO  ] Configuring VDSM
[ INFO  ] Starting vdsmd
[ INFO  ] Waiting for VDSM hardware info
[ INFO  ] Creating Storage Domain
[ INFO  ] Creating Storage Pool
[ INFO  ] Connecting Storage Pool
[ INFO  ] Verifying sanlock lockspace initialization
[ INFO  ] Initializing sanlock lockspace
[ INFO  ] Initializing sanlock metadata
[ INFO  ] Creating VM Image
[ INFO  ] Disconnecting Storage Pool
[ INFO  ] Start monitoring domain
[ INFO  ] Configuring VM
[ INFO  ] Updating hosted-engine configuration
[ INFO  ] Stage: Transaction commit
[ INFO  ] Stage: Closing up
[ INFO  ] Creating VM
You can now connect to the VM with the following command:
 /usr/bin/remote-viewer vnc://localhost:5900
Use temporary password "3042QHpX" to connect to vnc console.
Please note that in order to use remote-viewer you need to be able
to run graphical applications.
This means that if you are using ssh you have to supply the -Y flag
(enables trusted X11 forwarding).
Otherwise you can run the command from a terminal in your preferred
desktop environment.
If you cannot run graphical applications you can connect to the
graphic console from another host or connect to the console using
the following command:
virsh -c qemu+tls://Test/system console HostedEngine
If you need to reboot the VM you will need to start it manually
using the command:
hosted-engine --vm-start
You can then set a temporary password using the command:
hosted-engine --add-console-password
The VM has been started.  Install the OS and shut down or reboot
it.  To continue please make a selection:
```

```
          (1) Continue setup - VM installation is complete
          (2) Reboot the VM and restart installation
          (3) Abort setup

          (1, 2, 3)[1]:
```

Using the naming convention of this procedure, you would connect to the virtual machine using VNC with the following command:

```
/usr/bin/remote-viewer vnc://Host-HE1.example.com:5900
```

8. **Installing the Virtual Machine Operating System**

   Connect to HostedEngine-VM, the virtual machine created by the hosted-engine script, and install a Red Hat Enterprise Linux 6.5 operating system. Ensure the machine is rebooted once installation has completed.

9. **Synchronizing the Host and the Virtual Machine**

   Return to Host-HE1 and continue the **`hosted-engine`** deployment script by selecting option 1:

```
(1) Continue setup - VM installation is complete
```

```
 Waiting for VM to shut down...
[ INFO  ] Creating VM
You can now connect to the VM with the following command:
 /usr/bin/remote-viewer vnc://localhost:5900
Use temporary password "3042QHpX" to connect to vnc console.
Please note that in order to use remote-viewer you need to be able
to run graphical applications.
This means that if you are using ssh you have to supply the -Y flag
(enables trusted X11 forwarding).
Otherwise you can run the command from a terminal in your preferred
desktop environment.
If you cannot run graphical applications you can connect to the
graphic console from another host or connect to the console using
the following command:
virsh -c qemu+tls://Test/system console HostedEngine
If you need to reboot the VM you will need to start it manually
using the command:
hosted-engine --vm-start
You can then set a temporary password using the command:
hosted-engine --add-console-password
Please install and setup the engine in the VM.
You may also be interested in subscribing to "agent" RHN/Satellite
channel and installing rhevm-guest-agent-common package in the VM.
To continue make a selection from the options below:
          (1) Continue setup - engine installation is complete
          (2) Power off and restart the VM
          (3) Abort setup
```

10. **Installing the Manager**

Connect to HostedEngine-VM, subscribe to the appropriate Red Hat Enterprise Virtualization Manager channels, ensure that the most up-to-date versions of all installed packages are in use, and install the *rhevm* packages.

```
# yum upgrade
```

```
# yum install rhevm
```

11. **Configuring the Manager**

    Configure the engine on HostedEngine-VM:

    ```
    # engine-setup
    ```

12. **Synchronizing the Host and the Manager**

    Return to Host-HE1 and continue the **hosted-engine** deployment script by selecting option 1:

    ```
    (1) Continue setup - engine installation is complete
    ```

    ```
    [ INFO  ] Engine replied: DB Up!Welcome to Health Status!
    [ INFO  ] Waiting for the host to become operational in the engine.
    This may take several minutes...
    [ INFO  ] Still waiting for VDSM host to become operational...
    [ INFO  ] The VDSM Host is now operational
              Please shutdown the VM allowing the system to launch it
    as a monitored service.
              The system will wait until the VM is down.
    ```

13. **Shutting Down HostedEngine-VM**

    Shutdown HostedEngine-VM.

    ```
    # shutdown now
    ```

14. **Setup Confirmation**

    Return to Host-HE1 to confirm it has detected that HostedEngine-VM is down.

    ```
    [ INFO  ] Enabling and starting HA services
              Hosted Engine successfully set up
    [ INFO  ] Stage: Clean up
    [ INFO  ] Stage: Pre-termination
    [ INFO  ] Stage: Termination
    ```

Result

When the **hosted-engine** deployment script completes successfully, the Red Hat Enterprise Virtualization Manager is configured and running on your server. In contrast to a bare-metal Manager installation, the hosted engine Manager has already configured the data center, cluster, host (Host-HE1), storage domain, and virtual machine of the hosted engine (HostedEngine-VM). You

can log in as the **admin@internal** user to continue configuring the Manager and add further resources.

Link your Red Hat Enterprise Virtualization Manager to a directory server so you can add additional users to the environment. Red Hat Enterprise Virtualization supports directory services from Red Hat Directory Services (RHDS), IdM, and Active Directory. Add a directory server to your environment using the **engine-manage-domains** command.

The **ovirt-host-engine-setup** script also saves the answers you gave during configuration to a file, to help with disaster recovery. If a destination is not specified using the **--generate-answer=<file>** argument, the answer file is generated at **/etc/ovirt-hosted-engine/answers.conf**.

Report a bug

4.5. Migrating to a Self-Hosted Environment

Summary

Deploy a hosted-engine environment and migrate an existing instance of Red Hat Enterprise Virtualization. The **hosted-engine** deployment script is provided to assist with this task. The script asks you a series of questions, and configures your environment based on your answers. When the required values have been provided, the updated configuration is applied and the Red Hat Enterprise Virtualization Manager services are started.

The **hosted-engine** deployment script guides you through several distinct configuration stages. The script suggests possible configuration defaults in square brackets. Where these default values are acceptable, no additional input is required.

This procedure requires a new Red Hat Enterprise Linux 6.5 host with the *ovirt-hosted-engine-setup* package installed. This host is referred to as 'Host-HE1', with a fully qualified domain name (FQDN) of **Host-HE1.example.com** in this procedure.

Your original Red Hat Enterprise Virtualization Manager is referred to as 'BareMetal-Manager', with an FQDN of Host-HE1.example.com, in this procedure. You will be required to access and make changes on BareMetal-Manager during this procedure.

The hosted engine, the virtual machine created during configuration of Host-HE1 and used to manage the environment, is referred to as 'HostedEngine-VM' in this procedure. You will be prompted by the **hosted-engine** deployment script to access this virtual machine multiple times to install an operating system and to configure the engine.

All steps in this procedure are to be conducted as the **root** user for the specified machine.

Important

The engine running on BareMetal-Manager must be the same version as will be installed on HostedEngine-VM. As the hosted engine feature is only available on Red Hat Enterprise Virtualization version 3.3.0 and later, any previous version of Red Hat Enterprise Virtualization running on BareMetal-Manager must be upgraded. Upgrade the engine version on BareMetal-Manager before creating the backup with the **engine-backup** command.

Procedure 4.3. Migrating to a Self-Hosted Environment

1. **Initiating Hosted Engine Deployment**

Begin configuration of the self-hosted environment by deploying the **hosted-engine** customization script on Host_HE1. To escape the script at any time, use the **CTRL+D** keyboard combination to abort deployment.

```
# hosted-engine --deploy
```

2. **Configuring Storage**

   Select the version of NFS and specify the full address, using either the FQDN or IP address, and path name of the shared storage domain. Choose the storage domain and storage data center names to be used in the environment.

```
During customization use CTRL-D to abort.
Please specify the storage you would like to use (nfs3, nfs4)[nfs3]:
Please specify the full shared storage connection path to use
(example: host:/path): storage.example.com:/hosted_engine/nfs
[ INFO  ] Installing on first host
Please provide storage domain name. [hosted_storage]:
Local storage datacenter name is an internal name and currently
will not be shown in engine's admin UI.Please enter local
datacenter name [hosted_datacenter]:
```

3. **Configuring the Network**

   The script detects possible network interface controllers (NICs) to use as a management bridge for the environment. It then checks your firewall configuration and offers to modify it for console (SPICE or VNC) access HostedEngine-VM. Provide a pingable gateway IP address, to be used by the **ovirt-ha-agent** to help determine a host's suitability for running HostedEngine-VM.

```
Please indicate a nic to set rhevm bridge on: (eth1, eth0) [eth1]:
iptables was detected on your computer, do you wish setup to
configure it? (Yes, No)[Yes]:
Please indicate a pingable gateway IP address [X.X.X.X]:
```

4. **Configuring the Virtual Machine**

   The script creates a virtual machine to be configured as the Red Hat Enterprise Virtualization Manager, the hosted engine referred to in this procedure as HostedEngine-VM. Specify the boot device and, if applicable, the path name of the installation media, the CPU type, the number of virtual CPUs, and the disk size. Specify a MAC address for the HostedEngine-VM, or accept a randomly generated one. The MAC address can be used to update your DHCP server prior to installing the operating system on the virtual machine. Specify memory size and console connection type for the creation of HostedEngine-VM.

```
Please specify the device to boot the VM from (cdrom, disk, pxe)
[cdrom]:
The following CPU types are supported by this host:
          - model_Penryn: Intel Penryn Family
          - model_Conroe: Intel Conroe Family
Please specify the CPU type to be used by the VM [model_Penryn]:
Please specify the number of virtual CPUs for the VM [Defaults to
minimum requirement: 2]:
Please specify the disk size of the VM in GB [Defaults to minimum
```

```
requirement: 25]:
You may specify a MAC address for the VM or accept a randomly
generated default [00:16:3e:77:b2:a4]:
Please specify the memory size of the VM in MB [Defaults to minimum
requirement: 4096]:
Please specify the console type you would like to use to connect to
the VM (vnc, spice) [vnc]:
```

5. **Configuring the Hosted Engine**

Specify the name for Host-HE1 to be identified in the Red Hat Enterprise Virtualization environment, and the password for the **admin@internal** user to access the Administrator Portal. Provide the FQDN for HostedEngine-VM; this procedure uses the FQDN *HostedEngine-VM.example.com*. Finally, provide the name and TCP port number of the SMTP server, the email address used to send email notifications, and a comma-separated list of email addresses to receive these notifications.

Important

The FQDN provided for the engine (HostedEngine-VM.example.com) must be the same FQDN provided when BareMetal-Manager was initially set up.

```
Enter the name which will be used to identify this host inside the
Administrator Portal [hosted_engine_1]: Host-HE1
Enter 'admin@internal' user password that will be used for
accessing the Administrator Portal:
Confirm 'admin@internal' user password:
Please provide the FQDN for the engine you would like to use. This
needs to match the FQDN that you will use for the engine
installation within the VM: BareMetal-Manager.example.com
Please provide the name of the SMTP server through which we will
send notifications [localhost]:
Please provide the TCP port number of the SMTP server [25]:
Please provide the email address from which notifications will be
sent [root@localhost]:
Please provide a comma-separated list of email addresses which will
get notifications [root@localhost]:
```

6. **Configuration Preview**

Before proceeding, the **hosted-engine** script displays the configuration values you have entered, and prompts for confirmation to proceed with these values.

```
Bridge interface                   : eth1
Engine FQDN                        : BareMetal-Manager.example.com
Bridge name                        : rhevm
SSH daemon port                    : 22
Firewall manager                   : iptables
Gateway address                    : X.X.X.X
Host name for web application      : Host-HE1
Host ID                            : 1
Image size GB                      : 25
Storage connection                 :
```

```
storage.example.com:/hosted_engine/nfs
Console type                       : vnc
Memory size MB                     : 4096
MAC address                        : 00:16:3e:77:b2:a4
Boot type                          : pxe
Number of CPUs                     : 2
CPU Type                           : model_Penryn

Please confirm installation settings (Yes, No)[No]:
```

7. **Creating HostedEngine-VM**

   The script creates the virtual machine that will be configured to be HostedEngine-VM and provides connection details. You will need to install an operating system on HostedEngine-VM before the **hosted-engine** script can proceed on Host-HE1.

```
[ INFO  ] Generating answer file '/etc/ovirt-hosted-
engine/answers.conf'
[ INFO  ] Stage: Transaction setup
[ INFO  ] Stage: Misc configuration
[ INFO  ] Stage: Package installation
[ INFO  ] Stage: Misc configuration
[ INFO  ] Configuring libvirt
[ INFO  ] Generating VDSM certificates
[ INFO  ] Configuring VDSM
[ INFO  ] Starting vdsmd
[ INFO  ] Waiting for VDSM hardware info
[ INFO  ] Creating Storage Domain
[ INFO  ] Creating Storage Pool
[ INFO  ] Connecting Storage Pool
[ INFO  ] Verifying sanlock lockspace initialization
[ INFO  ] Initializing sanlock lockspace
[ INFO  ] Initializing sanlock metadata
[ INFO  ] Creating VM Image
[ INFO  ] Disconnecting Storage Pool
[ INFO  ] Start monitoring domain
[ INFO  ] Configuring VM
[ INFO  ] Updating hosted-engine configuration
[ INFO  ] Stage: Transaction commit
[ INFO  ] Stage: Closing up
[ INFO  ] Creating VM
You can now connect to the VM with the following command:
 /usr/bin/remote-viewer vnc://localhost:5900
Use temporary password "5379skAb" to connect to vnc console.
Please note that in order to use remote-viewer you need to be able
to run graphical applications.
This means that if you are using ssh you have to supply the -Y flag
(enables trusted X11 forwarding).
Otherwise you can run the command from a terminal in your preferred
desktop environment.
If you cannot run graphical applications you can connect to the
graphic console from another host or connect to the console using
the following command:
virsh -c qemu+tls://Test/system console HostedEngine
If you need to reboot the VM you will need to start it manually
using the command:
```

```
hosted-engine --vm-start
You can then set a temporary password using the command:
hosted-engine --add-console-password
The VM has been started.  Install the OS and shut down or reboot
it.  To continue please make a selection:

          (1) Continue setup - VM installation is complete
          (2) Reboot the VM and restart installation
          (3) Abort setup

          (1, 2, 3)[1]:
```

Using the naming convention of this procedure, you would connect to the virtual machine using VNC with the following command:

```
/usr/bin/remote-viewer vnc://Host-HE1.example.com:5900
```

8. **Installing the Virtual Machine Operating System**

   Connect to HostedEngine-VM, the virtual machine created by the hosted-engine script, and install a Red Hat Enterprise Linux 6.5 operating system.

9. **Synchronizing the Host and the Virtual Machine**

   Return to Host-HE1 and continue the **`hosted-engine`** deployment script by selecting option 1:

```
(1) Continue setup - VM installation is complete
```

```
 Waiting for VM to shut down...
[ INFO  ] Creating VM
You can now connect to the VM with the following command:
 /usr/bin/remote-viewer vnc://localhost:5900
Use temporary password "5379skAb" to connect to vnc console.
Please note that in order to use remote-viewer you need to be able
to run graphical applications.
This means that if you are using ssh you have to supply the -Y flag
(enables trusted X11 forwarding).
Otherwise you can run the command from a terminal in your preferred
desktop environment.
If you cannot run graphical applications you can connect to the
graphic console from another host or connect to the console using
the following command:
virsh -c qemu+tls://Test/system console HostedEngine
If you need to reboot the VM you will need to start it manually
using the command:
hosted-engine --vm-start
You can then set a temporary password using the command:
hosted-engine --add-console-password
Please install and setup the engine in the VM.
You may also be interested in subscribing to "agent" RHN/Satellite
channel and installing rhevm-guest-agent-common package in the VM.
```

```
To continue make a selection from the options below:
          (1) Continue setup - engine installation is complete
          (2) Power off and restart the VM
          (3) Abort setup
```

10. **Installing the Manager**

    Connect to HostedEngine-VM, subscribe to the appropriate Red Hat Enterprise Virtualization Manager channels, ensure that the most up-to-date versions of all installed packages are in use, and install the *rhevm* packages.

    ```
    # yum upgrade
    ```

    ```
    # yum install rhevm
    ```

11. **Disabling BareMetal-Manager**

    Connect to BareMetal-Manager, the Manager of your established Red Hat Enterprise Virtualization environment, and stop the engine and prevent it from running.

    ```
    # service ovirt-engine stop
    # service ovirt-engine disable
    # chkconfig ovirt-engine off
    ```

    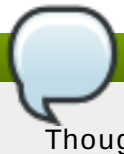

    Note

    Though stopping BareMetal-Manager from running is not obligatory, it is recommended as it ensures no changes will be made to the environment after the backup has been created. Additionally, it prevents BareMetal-Manager and HostedEngine-VM from simultaneously managing existing resources.

12. **Updating DNS**

    Update your DNS so that the FQDN of the Red Hat Enterprise Virtualization environment correlates to the IP address of HostedEngine-VM and the FQDN previously provided when configuring the **hosted-engine** deployment script on Host-HE1. In this procedure that FQDN was set as *BareMetal-Manager.example.com* because in a migrated hosted-engine setup, the FQDN provided for the engine must be identical to that given in the engine setup of the original engine.

13. **Creating a Backup of BareMetal-Manager**

    Connect to BareMetal-Manager and run the **engine-backup** command with the ***--mode=backup***, ***--file=FILE***, and ***--log=LogFILE*** parameters to specify the backup mode, the name of the backup file created and used for the backup, and the name of the log file to be created to store the backup log.

    ```
    # engine-backup --mode=backup --file=FILE --log=LogFILE
    ```

14. **Copying the Backup File to HostedEngine-VM**

Still on BareMetal-Manager, secure copy the backup file to HostedEngine-VM. In the following example, *HostedEngine-VM.example.com* is the FQDN for HostedEngine-VM, and */backup/* is any designated folder or path. If the designated folder or path does not exist, you will need to connect to HostedEngine-VM and create it before secure copying the backup from BareMetal-Manager.

```
# scp -p backup1 HostedEngine-VM.example.com:/backup/
```

15. **Restoring the Backup File on HostedEngine-VM**

The **engine-backup --mode=restore** command does not create a database; you are required to create one on HostedEngine-VM before restoring the backup you created on BareMetal-Manager. Connect to HostedEngine-VM and create the database, as detailed in Section 3.7, “Preparing a PostgreSQL Database for Use with Red Hat Enterprise Virtualization Manager”.

Note

The procedure in Section 3.7, “Preparing a PostgreSQL Database for Use with Red Hat Enterprise Virtualization Manager” creates a database that is not empty, which will result in the following error when you attempt to restore the backup:

```
FATAL: Database is not empty
```

Create an empty database using the following command in psql:

```
postgres=# create database [database name] owner [user name]
```

After the empty database has been created, restore the BareMetal-Manager backup using the **engine-backup** command with the ***--mode=restore --file=FILE --log=Restore.log*** parameters to specify the restore mode, the name of the file to be used to restore the database, and the name of the logfile to store the restore log. This restores the files and the database but does not start the service.

To specify a different database configuration, use the ***--change-db-credentials*** parameter to activate alternate credentials. Use the **engine-backup --help** command on the Manager for a list of credential parameters.

```
# engine-backup --mode=restore --file=FILE --log=Restore.log --
change-db-credentials --db-host=X.X.X.X --db-user=engine --db-
password=password --db-name=engine
```

16. **Configuring HostedEngine-VM**

Configure the engine on HostedEngine-VM. This will identify the existing files and database.

```
# engine-setup
```

```
[ INFO  ] Stage: Initializing
[ INFO  ] Stage: Environment setup
```

```
Configuration files: ['/etc/ovirt-engine-setup.conf.d/10-
packaging.conf', '/etc/ovirt-engine-setup.conf.d/20-setup-ovirt-
post.conf']
Log file: /var/log/ovirt-engine/setup/ovirt-engine-setup-
20140304075238.log
Version: otopi-1.1.2 (otopi-1.1.2-1.el6ev)
[ INFO  ] Stage: Environment packages setup
[ INFO  ] Yum Downloading: rhel-65-zstream/primary_db 2.8 M(70%)
[ INFO  ] Stage: Programs detection
[ INFO  ] Stage: Environment setup
[ INFO  ] Stage: Environment customization

          --== PACKAGES ==--

[ INFO  ] Checking for product updates...
[ INFO  ] No product updates found

          --== NETWORK CONFIGURATION ==--

Setup can automatically configure the firewall on this system.
Note: automatic configuration of the firewall may overwrite current
settings.
Do you want Setup to configure the firewall? (Yes, No) [Yes]:
[ INFO  ] iptables will be configured as firewall manager.

          --== DATABASE CONFIGURATION ==--


          --== OVIRT ENGINE CONFIGURATION ==--

          Skipping storing options as database already prepared

          --== PKI CONFIGURATION ==--

          PKI is already configured

          --== APACHE CONFIGURATION ==--


          --== SYSTEM CONFIGURATION ==--


          --== END OF CONFIGURATION ==--

[ INFO  ] Stage: Setup validation
[WARNING] Less than 16384MB of memory is available
[ INFO  ] Cleaning stale zombie tasks

          --== CONFIGURATION PREVIEW ==--

          Database name                    : engine
          Database secured connection      : False
          Database host                    : X.X.X.X
          Database user name               : engine
          Database host name validation    : False
          Database port                    : 5432
```

```
          NFS setup                                : True
          Firewall manager                         : iptables
          Update Firewall                          : True
          Configure WebSocket Proxy                : True
          Host FQDN                                : HostedEngine-
VM.example.com
          NFS mount point                          : /var/lib/exports/iso
          Set application as default page          : True
          Configure Apache SSL                     : True

          Please confirm installation settings (OK, Cancel) [OK]:
```

Confirm the settings. Upon completion, the setup will provide an SSH fingerprint and an internal Certificate Authority hash.

17. **Synchronizing the Host and the Manager**

    Return to Host-HE1 and continue the **hosted-engine** deployment script by selecting option 1:

```
(1) Continue setup - engine installation is complete
```

```
[ INFO  ] Engine replied: DB Up!Welcome to Health Status!
[ INFO  ] Waiting for the host to become operational in the engine.
This may take several minutes...
[ INFO  ] Still waiting for VDSM host to become operational...
[ INFO  ] The VDSM Host is now operational
          Please shutdown the VM allowing the system to launch it
as a monitored service.
          The system will wait until the VM is down.
```

18. **Shutting Down HostedEngine-VM**

    Shutdown HostedEngine-VM.

```
# shutdown now
```

19. **Setup Confirmation**

    Return to Host-HE1 to confirm it has detected that HostedEngine-VM is down.

```
[ INFO  ] Enabling and starting HA services
          Hosted Engine successfully set up
[ INFO  ] Stage: Clean up
[ INFO  ] Stage: Pre-termination
[ INFO  ] Stage: Termination
```

Result

Your Red Hat Enterprise Virtualization engine has been migrated to a hosted-engine setup. The Manager is now operating on a virtual machine on Host-HE1, called HostedEngine-VM in the environment. As HostedEngine-VM is highly available, it will be migrated to other hosts in the environment when applicable.

Report a bug

4.6. Installing Additional Hosts to a Self-Hosted Environment

Summary

Adding additional nodes to a self-hosted environment is very similar to deploying the original host, though heavily truncated as the script detects the environment.

As with the original host, additional hosts require Red Hat Enterprise Linux 6.5 with subscriptions to the appropriate Red Hat Enterprise Virtualization channels.

All steps in this procedure are to be conducted as the **root** user.

Procedure 4.4. Adding the host

1. Install the *ovirt-hosted-engine-setup* package.

```
# yum install ovirt-hosted-engine-setup
```

2. Configure the host with the deployment command.

```
# hosted-engine --deploy
```

3. **Configuring Storage**

   Specify the storage type and the full address, using either the Fully Qualified Domain Name (FQDN) or IP address, and path name of the shared storage domain used in the self-hosted environment.

```
Please specify the storage you would like to use (nfs3, nfs4)[nfs3]:
Please specify the full shared storage connection path to use
(example: host:/path): storage.example.com:/hosted_engine/nfs
```

4. **Detecting the Self-Hosted Engine**

   The **hosted-engine** script detects that the shared storage is being used and asks if this is an additional host setup. You are then prompted for the host ID, which must be an integer not already assigned to an additional host in the environment.

```
The specified storage location already contains a data domain. Is
this an additional host setup (Yes, No)[Yes]?
[ INFO  ] Installing on additional host
Please specify the Host ID [Must be integer, default: 2]:
```

5. **Configuring the System**

   The **hosted-engine** script uses the answer file generated by the original hosted-engine setup. To achieve this, the script requires the FQDN or IP address and the password of the **root** user of that host so as to access and secure-copy the answer file to the additional host.

```
[WARNING] A configuration file must be supplied to deploy Hosted
Engine on an additional host.
The answer file may be fetched from the first host using scp.
```

```
If you do not want to download it automatically you can abort the
setup answering no to the following question.
Do you want to scp the answer file from the first host? (Yes, No)
[Yes]:
Please provide the FQDN or IP of the first host:
Enter 'root' user password for host Host-HE1.example.com:
[ INFO  ] Answer file successfully downloaded
```

6. **Configuring the Hosted Engine**

   Specify the name for the additional host to be identified in the Red Hat Enterprise Virtualization environment, and the password for the **admin@internal** user.

```
Enter the name which will be used to identify this host inside the
Administrator Portal [hosted_engine_2]:
Enter 'admin@internal' user password that will be used for
accessing the Administrator Portal:
Confirm 'admin@internal' user password:
```

7. **Configuration Preview**

   Before proceeding, the **hosted-engine** script displays the configuration values you have entered, and prompts for confirmation to proceed with these values.

```
Bridge interface                   : eth1
Engine FQDN                        : HostedEngine-VM.example.com
Bridge name                        : rhevm
SSH daemon port                    : 22
Firewall manager                   : iptables
Gateway address                    : X.X.X.X
Host name for web application      : hosted_engine_2
Host ID                            : 2
Image size GB                      : 25
Storage connection                 :
storage.example.com:/hosted_engine/nfs
Console type                       : vnc
Memory size MB                     : 4096
MAC address                        : 00:16:3e:05:95:50
Boot type                          : disk
Number of CPUs                     : 2
CPU Type                           : model_Penryn

Please confirm installation settings (Yes, No)[No]:
```

Result

After confirmation, the script completes installation of the host and adds it to the environment.

Report a bug

4.7. Maintaining the Self-Hosted Engine

Setting the engine to global maintenance enables you to start, stop, and modify the engine without interference from the high availability agents. This must be applied to the engine for any setup or upgrade operations that require the engine to be stopped, for instance the installation of the *rhevm-dwh* and *rhevm-reports* packages necessary for the **Reports Portal**.

```
# hosted-engine --set-maintenance --mode=global
```

To resume the high availability function of the engine, turn off global maintenance:

```
# hosted-engine --set-maintenance --mode=none
```

Both of these commands are to be conducted as the **root** user.

Report a bug

Chapter 5. Data Collection Setup and Reports Installation

5.1. Workflow Progress — Data Collection Setup and Reports Installation

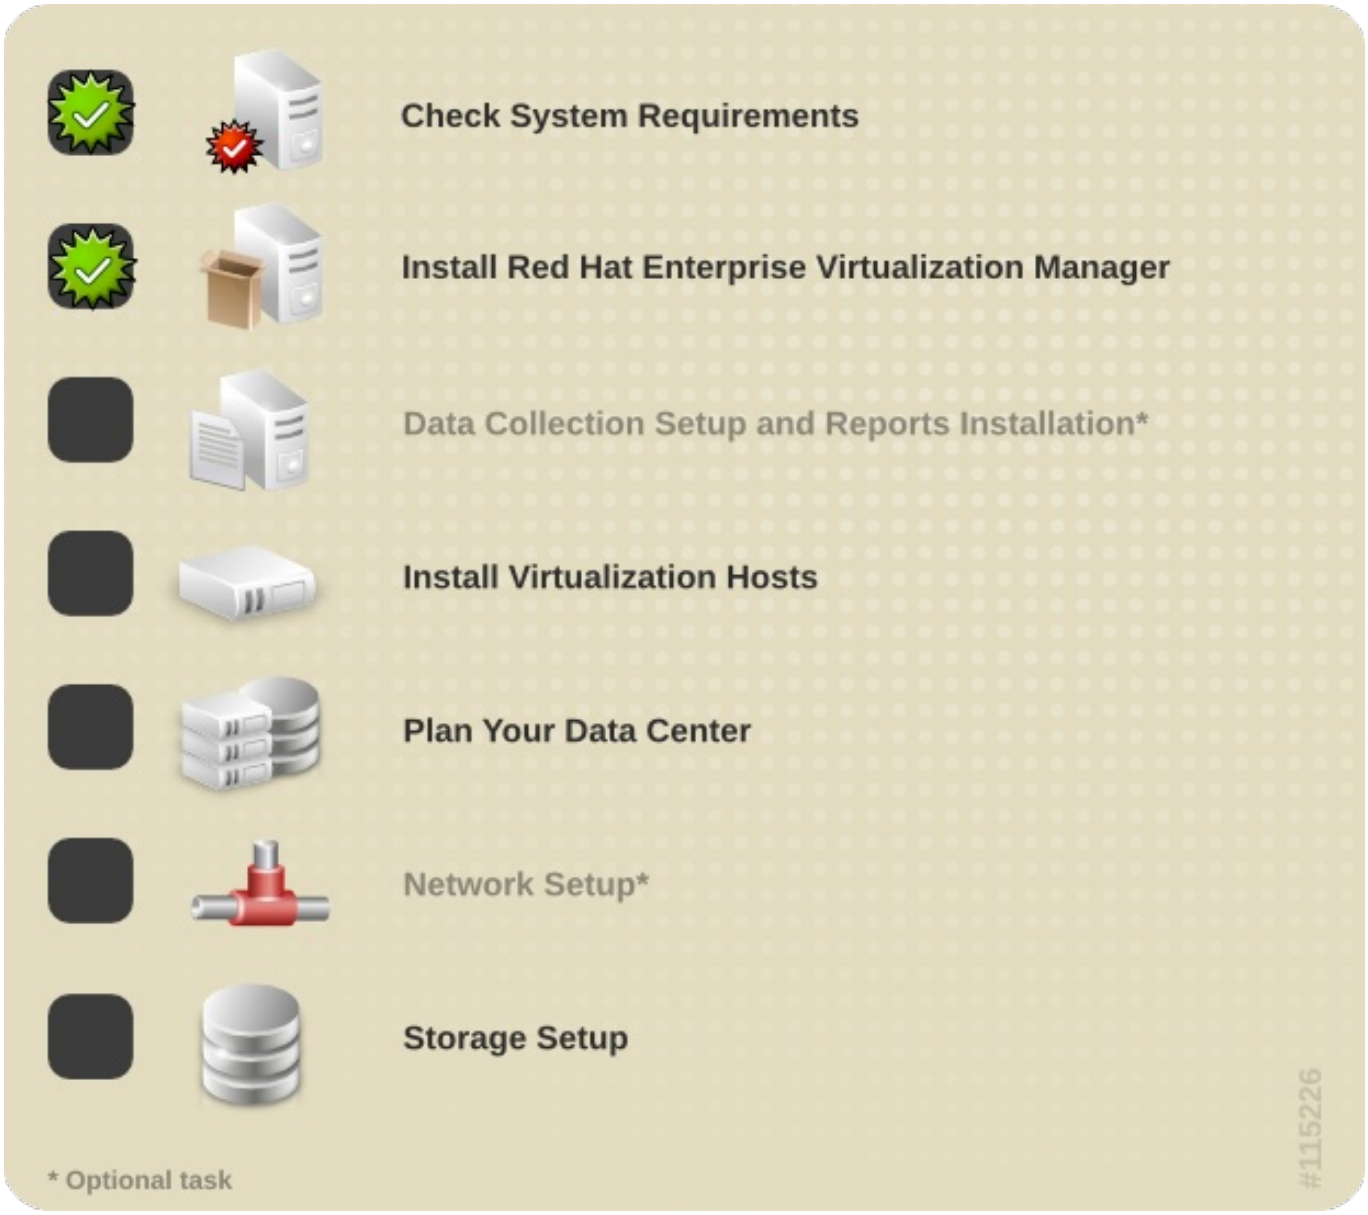

Report a bug

5.2. Data Collection Setup and Reports Installation Overview

The Red Hat Enterprise Virtualization Manager optionally includes a comprehensive management history database, which can be utilized by any application to extract a range of information at the data center, cluster, and host levels. As the database structure changes over time a number of database views are also included to provide a consistent structure to consuming applications. A view is a virtual table composed of the result set of a database query. The definition of a view is stored in the database as a **`SELECT`** statement. The result set of the **`SELECT`** statement populates the virtual table returned by the view. If the optional comprehensive management history database has been enabled, the history tables and their associated views are stored in the **`ovirt_engine_history`** database.

In addition to the history database Red Hat Enterprise Virtualization Manager Reports functionality is also available as an optional component. Red Hat Enterprise Virtualization Manager Reports provides a customized implementation of JasperServer, and JasperReports. JasperServer is a component of JasperReports, an open source reporting tool capable of being embedded in Java-based applications. It produces reports which can be rendered to screen, printed, or exported to a variety of formats including PDF, Excel, CSV, Word, RTF, Flash, ODT and ODS. Reports built in Red Hat Enterprise Virtualization Manager Reports are accessed via a web interface. In addition to a range of pre-configured reports and dashboards for monitoring the system, you are also able to create your own ad hoc reports.

Before proceeding with Red Hat Virtualization Manager Reports installation you must first have installed the Red Hat Enterprise Virtualization Manager.

The Red Hat Enterprise Virtualization Manager Reports functionality depends on the presence of the history database, which is installed separately. Both the history database and the Red Hat Enterprise Virtualization Manager Reports are optional components. They are not installed by default when you install the Red Hat Enterprise Virtualization Manager.

Report a bug

5.3. Installing and Configuring the History Database

Prerequisites:

- Section 3.4, “Installing the Red Hat Enterprise Virtualization Manager Packages”
- Section 3.5, “Configuring Red Hat Enterprise Virtualization Manager”

Summary

Use of the history database is optional. To use the history or reporting capabilities of Red Hat Enterprise Virtualization Manager, you must install and configure the history database.

Procedure 5.1. Installing and Configuring the History Database

1. Install the *rhevm-dwh* package. This package must be installed on the system on which the Red Hat Enterprise Virtualization Manager is installed.

   ```
   # yum install rhevm-dwh
   ```

2. Once the required packages have been downloaded, they are listed for review. You will be prompted to confirm continuing with the installation. Upon confirmation, the packages are installed. However, some further configuration is required before the reports functionality can be used.

3. Configure the history database. Use the **rhevm-dwh-setup** command to configure the Extract, Transform, Load (ETL) process and database scripts used to create and maintain a working history database.

   a. Run the **rhevm-dwh-setup** command on the system hosting the Red Hat Enterprise Virtualization Manager:

      ```
      # rhevm-dwh-setup
      ```

   b. For the history database installation to take effect, the **ovirt-engine** service must be restarted. The **rhevm-dwh-setup** command prompts you:

```
In order to proceed the installer must stop the ovirt-engine
service
Would you like to stop the ovirt-engine service? (yes|no):
```

Type **yes** and then press **Enter** to proceed.

c. The **rhevm-dwh-setup** utility can optionally create a read-only user to allow remote access to the history database.

```
This utility can configure a read only user for DB access.
Would you like to do so? (yes|no):
Provide a username for read-only user:
Provide a password for read-only user:
```

If you choose to create a read-only user, the **rhevm-dwh-setup** utility automatically opens the required firewall ports and configures the database to listen on externally facing network interface devices.

Note

The **rhevm-dwh-setup** utility can configure read-only access to **local** databases only. If you chose to use a remote database during **engine-setup**, you have to manually configure read-only access to the history database. See *Connecting to the History Database* in the *Red Hat Enterprise Virtualization Administration Guide*.

d. The **rhevm-dwh-setup** utility can optionally configure the history database to use secure connections

```
Should postgresql be setup with secure connection? (yes|no):
```

The command then creates and configures the **ovirt_engine_history** database and starts the **ovirt-engine** service.

Result

The **ovirt_engine_history** database has been created. Red Hat Enterprise Virtualization Manager is configured to log information to this database for reporting purposes.

Report a bug

5.4. Installing and Configuring Red Hat Enterprise Virtualization Manager Reports

Prerequisites:

- Section 5.3, “Installing and Configuring the History Database”

Summary

Use of reports is optional. To use the reporting capabilities of Red Hat Enterprise Virtualization Manager, you must install and configure *rhevm-reports*.

Procedure 5.2. Installing and Configuring Red Hat Enterprise Virtualization Manager Reports

1. Install the *rhevm-reports* package. This package must be installed on the system on which the Red Hat Enterprise Virtualization Manager is installed.

```
# yum install rhevm-reports
```

2. Run the **rhevm-reports-setup** command on the system hosting the Red Hat Enterprise Virtualization Manager:

```
# rhevm-reports-setup
```

3. For the Red Hat Enterprise Virtualization Manager Reports installation to take effect, the **ovirt-engine** service must be restarted. The **rhevm-reports-setup** command prompts you:

```
In order to proceed the installer must stop the ovirt-engine
service
Would you like to stop the ovirt-engine service? (yes|no):
```

Type **yes** and then press **Enter** to proceed. The command then performs a number of actions before prompting you to set the password for the Red Hat Enterprise Virtualization Manager Reports administrative users (**rhevm-admin** and **superuser**). Note that the reports system maintains its own set of credentials which are separate to those used for Red Hat Enterprise Virtualization Manager.

```
Please choose a password for the reports admin user(s) (rhevm-admin
and superuser):
```

You will be prompted to enter the password a second time to confirm it.

Result

Red Hat Enterprise Virtualization Manager Reports has been installed successfully. Access Red Hat Enterprise Virtualization Manager Reports at **http://*[demo.redhat.com]*/rhevm-reports**, replacing ***[demo.redhat.com]*** with the fully-qualified domain name of the Red Hat Enterprise Virtualization Manager. If during Red Hat Enterprise Virtualization Manager installation you selected a non-default HTTP port then append **:***[port]* to the URL, replacing *[port]* with the port that you chose.

Use the user name **rhevm-admin** and the password you set during reports installation to log in for the first time. Note that the first time you log into Red Hat Enterprise Virtualization Manager Reports, a number of web pages are generated, and as a result your initial attempt to login may take some time to complete.

Report a bug

Chapter 6. Updating the Red Hat Enterprise Virtualization Environment

6.1. Upgrades between Minor Releases

6.1.1. Checking for Red Hat Enterprise Virtualization Manager Updates

Summary

Use the **engine-upgrade-check** command, included in Red Hat Enterprise Virtualization Manager, to check for updates.

Procedure 6.1. Checking for Red Hat Enterprise Virtualization Manager Updates

1. Run **engine-upgrade-check** as a user with administrative privileges such as the **root** user.

```
# engine-upgrade-check
```

2. A. Where no updates are available the command will output the text **No upgrade**.

```
# engine-upgrade-check
VERB: queue package rhevm-setup for update
VERB: package rhevm-setup queued
VERB: Building transaction
VERB: Empty transaction
VERB: Transaction Summary:
No upgrade
```

   B. Where updates are available the command will list the packages to be updated.

```
# engine-upgrade-check
VERB: queue package rhevm-setup for update
VERB: package rhevm-setup queued
VERB: Building transaction
VERB: Transaction built
VERB: Transaction Summary:
VERB:     updated    - rhevm-lib-3.3.0-0.46.el6ev.noarch
VERB:     update     - rhevm-lib-3.3.1-0.48.el6ev.noarch
VERB:     updated    - rhevm-setup-3.3.0-0.46.el6ev.noarch
VERB:     update     - rhevm-setup-3.3.1-0.48.el6ev.noarch
Upgrade available
```

Result

You have successfully checked for Red Hat Enterprise Virtualization Manager updates.

Report a bug

6.1.2. Updating Red Hat Enterprise Virtualization Manager

Summary

Updates to the Red Hat Enterprise Virtualization Manager are released via Red Hat Network. Before installing an update from Red Hat Network be sure to read the advisory text associated with it as well as the most recent version of the *Red Hat Enterprise Virtualization Release Notes*, and the *Red Hat Enterprise Virtualization Technical Notes*. A number of actions must be performed to complete an upgrade, including:

- Stopping the **ovirt-engine** service.
- Downloading and installing the updated packages.
- Backing up and updating the database.
- Performing post installation configuration.
- Restarting the **ovirt-engine** service.

A script is included to perform these actions for you in an automated fashion. Active virtualization hosts are not upgraded by this process and must be updated separately. As a result the virtual machines running upon them are not affected. All commands in this task must be run while logged into the system hosting Red Hat Enterprise Virtualization Manager as the **root** user.

Procedure 6.2. Updating Red Hat Enterprise Virtualization Manager

1. Run the **yum** command to update the *rhevm-setup* package.

```
# yum update rhevm-setup
```

2. Run the **engine-setup** command to update the Red Hat Enterprise Virtualization Manager.

```
# engine-setup
```

Note

From Version 3.3, installation of Red Hat Enterprise Virtualization Manager supports **otopi**, a standalone, plug-in-based installation framework for setting up system components. Under this framework, the **rhevm-upgrade** command used during the installation process has been updated to **engine-setup** and is now obsolete.

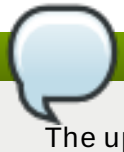

Note

The upgrade process may take some time; allow time for the upgrade process to complete and do not stop the process once initiated. Once the upgrade has been completed, you will also be instructed to separately upgrade the data warehouse and reports functionality. These additional steps are only required if these optional packages are installed.

Result

You have successfully updated Red Hat Enterprise Virtualization Manager.

Report a bug

6.1.3. Troubleshooting for Upgrading Red Hat Enterprise Virtualization Manager

Red Hat Enterprise Virtualization Troubleshooting Cases

SAM Channel Causes Conflicts with rhevm upgrade

Running Red Hat Enterprise Virtualization Manager on a machine that has Subscription Asset Manager (SAM) enabled is not supported. The **yum update** command fails to update **rhevm** due to a "file conflicts" error if the **sam-rhel-x86_64-server-6** channel is enabled.

If your Red Hat Enterprise Virtualization environment does not require Subscription Asset Manager (SAM) features, you can disable the following channels in the customer portal:

1. **sam-rhel-x86_64-server-6**
2. **sam-rhel-x86_64-server-6-debuginfo**

Then, remove the package causing the conflict by issuing this command:

```
# yum remove apache-commons-codec
```

Alternatively, remove the channels from the command line:

```
# rhn-channel -r -c sam-rhel-x86_64-server-6
# rhn-channel -r -c sam-rhel-x86_64-server-6-debuginfo
```

Then, remove the package causing the conflict by issuing this command:

```
# yum remove apache-commons-codec
```

Report a bug

6.1.4. Updating Red Hat Enterprise Virtualization Manager Reports

Summary

The Red Hat Enterprise Virtualization Manager Reports must be updated separately to the Red Hat Enterprise Virtualization Manager. This task provides instructions for updating both the history database and the reporting engine. You must run all commands in this task while logged into the system hosting Red Hat Enterprise Virtualization Manager as the **root** user.

Procedure 6.3. Updating Red Hat Enterprise Virtualization Manager Reports

1. Use the **yum** command to update the *rhevm-reports* and *rhevm-dwh* packages.

```
# yum update -y rhevm-reports rhevm-dwh
```

2. Run the **rhevm-dwh-setup** command to update the **ovirt_engine_history** database.

```
# rhevm-dwh-setup
```

3. Run the **rhevm-reports-setup** command to update the reporting engine.

```
# rhevm-reports-setup
```

Result

You have successfully updated the Red Hat Enterprise Virtualization Manager Reports.

Report a bug

6.1.5. Updating Red Hat Enterprise Virtualization Hypervisors

Summary

Updating Red Hat Enterprise Virtualization Hypervisors involves reinstalling the Hypervisor with a newer version of the Hypervisor ISO image. This includes stopping and restarting the Hypervisor. Virtual machines are automatically migrated to a different host, as a result it is recommended that Hypervisor updates are performed at a time when the host's usage is relatively low.

It is recommended that administrators update Red Hat Enterprise Virtualization Hypervisors regularly. Important bug fixes and security updates are included in updates. Hypervisors which are not up to date may be a security risk.

Warning

Upgrading Hypervisor hosts involves shutting down, deactivating guests, and restarting the physical server. If any virtual machines are running on the Hypervisor, all data and configuration details may be destroyed if they are not shut down. Upgrading Hypervisors must be carefully planned and executed with care and consideration.

Important

Ensure that the cluster contains more than one host before performing an upgrade. Do not attempt to re-install or upgrade all the hosts at the same time, as one host must remain available to perform Storage Pool Manager (SPM) tasks.

Procedure 6.4. Updating Red Hat Enterprise Virtualization Hypervisors

1. Log in to the system hosting Red Hat Enterprise Virtualization Manager as the **root** user.
2. Ensure that:
   - the system is subscribed to the **Red Hat Enterprise Virtualization** entitlement — if using certificate-based Red Hat Network; or
   - the system is subscribed to the **Red Hat Enterprise Virtualization Hypervisor (v.6 x86-64)** (labeled **rhel-x86_64-server-6-rhevh**) — if using classic Red Hat Network.
3. Run the **yum** command with the ***update rhev-hypervisor6*** parameters to ensure that you have the most recent version of the *rhev-hypervisor6* package installed.

```
# yum update rhev-hypervisor6
```

4. Use your web browser to log in to the Administration Portal as a Red Hat Enterprise Virtualization administrative user.

5. Click the **Hosts** tab, and then select the host that you intend to upgrade. If the host is not displayed, or the list of hosts is too long to filter visually, perform a search to locate the host.

6. With the host selected, click the **General** tab on the **Details** pane.

   A. If the host requires updating, an alert message indicates that a new version of the Red Hat Enterprise Virtualization Hypervisor is available.

   B. If the host does not require updating, no alert message is displayed and no further action is required.

7. Ensure the host remains selected and click the **Maintenance** button, if the host is not already in maintenance mode. This will cause any virtual machines running on the host to be migrated to other hosts. If the host is the SPM, this function will be moved to another host. The status of the host changes as it enters maintenance mode. When the host status is **Maintenance**, the message in the general tab changes, providing you with a link which when clicked will re-install or upgrade the host.

8. Ensure that the host remains selected, and that you are on the **General** tab of the the **Details** pane. Click the **Upgrade** link. The **Install Host** dialog box displays.

9. Select **rhev-hypervisor.iso**, which is symbolically linked to the most recent hypervisor image.

10. Click **OK** to update and re-install the host. The dialog closes, the details of the host are updated in the **Hosts** tab, and the status changes.

    The host status will transition through these stages:

    - **Installing**,
    - **Reboot**,
    - **Non Responsive**, and
    - **Up**.

    These are all expected, and each stage will take some time.

11. Once successfully updated, the host displays a status of **Up**. Any virtual machines that were migrated off the host, are at this point able to be migrated back to it.

Important

After a Red Hat Enterprise Virtualization Hypervisor is successfully registered to the Red Hat Enterprise Virtualization Manager and then upgraded, it may erroneously appear in the Administration Portal with the status of **Install Failed**. Click on the **Activate** button, and the hypervisor will change to an **Up** status and be ready for use.

Result

You have successfully updated a Red Hat Enterprise Virtualization Hypervisor. Repeat these steps for each Hypervisor in the Red Hat Enterprise Virtualization environment.

Report a bug

6.1.6. Updating Red Hat Enterprise Linux Virtualization Hosts

Summary

Red Hat Enterprise Linux virtualization hosts are updated the same way as regular Red Hat Enterprise Linux systems, using **yum**. It is highly recommended that you use **yum** to update your systems regularly, to ensure timely application of security and bug fixes. All steps in this task must be run while logged into the Red Hat Enterprise Linux virtualization host as the **root** user.

Procedure 6.5. Updating Red Hat Enterprise Linux Virtualization Hosts

1. On the administration portal, navigate to the **Hosts** tab and select the host to be updated. Click **Maintenance** to place it into maintenance mode.
2. On the Red Hat Enterprise Linux virtualization host, run the **yum** command with the ***update*** parameter to update all installed packages.

```
# yum update
```

3. If a package such as the kernel was updated, you must reboot the host to get the new functionality. If a package such as *VDSM* or *libvirt* was updated, you must restart that service to get the new functionality. Moreover, if the *libvirt* package is updated, you must restart the **VDSM** service.

Result

You have successfully updated the Red Hat Enterprise Linux virtualization host. Repeat this process for each Red Hat Enterprise Linux virtualization host in the Red Hat Enterprise Virtualization environment.

Report a bug

6.1.7. Updating the Red Hat Enterprise Virtualization Guest Tools

Summary

The guest tools allow the Manager to communicate with running guests, providing information such as IP addresses, memory usage, and installed applications for display in the Administration Portal. The guest tools are distributed as an ISO file that can be attached to guests. The ISO file is packaged as an RPM file, that can be installed and upgraded on registered Red Hat Enterprise Virtualization Manager servers.

Procedure 6.6. Updating the Red Hat Enterprise Virtualization Guest Tools

1. On the Manager, as root user, use the **yum upgrade** to upgrade the **rhev-guest-tools-iso** package.

```
# yum update -y rhev-guest-tools-iso*
```

2. When the **rhev-guest-tools-iso** package has been successfully upgraded, use the **engine-iso-uploader** command to upload it to your ISO storage domain. Replace *[ISODomain]* with the name of your ISO storage domain.

```
engine-iso-uploader --iso-domain=[ISODomain] upload
/usr/share/rhev-guest-tools-iso/rhev-tools-setup.iso
```

The **rhev-tools-setup.iso** file is actually a link to the most recently updated ISO file. The link is automatically changed to point to the newest ISO file every time you upgrade the **rhev-guest-tools-iso** package.

3. Using the web portal or REST API, attach the rhev-tools-setup.iso file to each of your guests, and from within each guest, upgrade the tools installed on each guest using the installer on the ISO.

Result

You've updated your **rhev-tools-setup.iso** file, uploaded the updated ISO to your ISO storage domain, and attached it to your virtual machines.

Report a bug

6.2. Upgrading to Red Hat Enterprise Virtualization 3.3

6.2.1. Red Hat Enterprise Virtualization Manager 3.3 Upgrade Overview

Upgrading Red Hat Enterprise Virtualization Manager is a straightforward process that comprises three main steps:

- Configuring channels and entitlements.
- Updating the required packages.
- Performing the upgrade.

The command used to perform the upgrade itself is **engine-setup**, which provides an interactive interface. While the upgrade is in process, virtualization hosts and the virtual machines running on those virtualization hosts continue to operate independently. When the upgrade is complete, you can then upgrade your hosts to the latest versions of Red Hat Enterprise Linux or Red Hat Enterprise Virtualization Hypervisor.

Report a bug

6.2.2. Red Hat Enterprise Virtualization 3.3 Upgrade Considerations

The following is a list of key considerations that must be made when planning your upgrade.

Important

Upgrading to version 3.3 can only be performed from version 3.2

Users of Red Hat Enterprise Virtualization 3.1 must migrate to Red Hat Enterprise Virtualization 3.2 before attempting to upgrade to Red Hat Enterprise Virtualization 3.3.

Red Hat Enterprise Virtualization Manager cannot be installed on the same machine as IPA

An error message displays if the *ipa-server* package is installed. Red Hat Enterprise Virtualization Manager 3.3 does not support installation on the same machine as Identity Management (IdM). To resolve this issue, you must migrate the IdM configuration to another system before re-attempting the upgrade. For further information, see https://access.redhat.com/knowledge/articles/233143.

```
Error: IPA was found to be installed on this machine. Red
Hat Enterprise Virtualization Manager 3.3 does not support
installing IPA on the same machine. Please remove ipa
packages before you continue.
```

Upgrading to JBoss Enterprise Application Platform 6.1.0 is recommended

Although Red Hat Enterprise Virtualization Manager 3.3 supports Enterprise Application Platform 6.0.1, upgrading to the latest supported version of JBoss is recommended. For more information on upgrading to JBoss Enterprise Application Platform 6.1.0, see Upgrade the JBoss EAP 6 RPM Installation.

The rhevm-upgrade command has been replaced by engine-setup

From Version 3.3, installation of Red Hat Enterprise Virtualization Manager supports **otopi**, a standalone, plug-in-based installation framework for setting up system components. Under this framework, the **rhevm-upgrade** command used during the installation process has been updated to **engine-setup** and is now obsolete.

Report a bug

6.2.3. Upgrading to Red Hat Enterprise Virtualization Manager 3.3

Prerequisites:

- Section 6.3.1, “Upgrading to Red Hat Enterprise Virtualization Manager 3.2”

Summary

The following procedure outlines the process for upgrading Red Hat Enterprise Virtualization Manager 3.2 to Red Hat Enterprise Virtualization Manager 3.3. This procedure assumes that the system on which the Manager is hosted is subscribed to the channels and entitlements for receiving Red Hat Enterprise Virtualization 3.2 packages.

If the upgrade fails, the **engine-setup** command will attempt to roll your Red Hat Enterprise Virtualization Manager installation back to its previous state. For this reason, the channels required by Red Hat Enterprise Virtualization 3.2 must not be removed until after the upgrade is complete as outlined below. If the upgrade fails, detailed instructions display that explain how to restore your installation.

Procedure 6.7. Upgrading to Red Hat Enterprise Virtualization Manager 3.3

1. Subscribe the system to the required channels and entitlements for receiving Red Hat Enterprise Virtualization Manager 3.3 packages.

   ### Subscription Manager

   Red Hat Enterprise Virtualization 3.3 packages are provided by the **rhel-6-server-rhevm-3.3-rpms** repository associated with the **Red Hat Enterprise Virtualization** entitlement. Use the **yum-config-manager** command to enable the repository in your **yum** configuration.

   ```
   # yum-config-manager --enable rhel-6-server-rhevm-3.3-rpms
   ```

   ### Red Hat Network Classic

   The Red Hat Enterprise Virtualization 3.3 packages are provided by the **Red Hat Enterprise Virtualization Manager (v.3.3 x86_64)** channel, also referred to as **rhel-x86_64-server-6-rhevm-3.3** in Red Hat Network Classic. Use the **rhn-channel** command or the Red Hat Network web interface to subscribe to the **Red Hat Enterprise Virtualization Manager (v.3.3 x86_64)** channel:

   ```
   # rhn-channel --add --channel=rhel-x86_64-server-6-rhevm-3.3
   ```

2. Update the *rhevm-setup* package to ensure you have the most recent version of **engine-setup**.

   ```
   # yum update rhevm-setup
   ```

3. Run the **engine-setup** command and follow the prompts to upgrade Red Hat Enterprise Virtualization Manager.

   ```
   # engine-setup
   [ INFO  ] Stage: Initializing

             Welcome to the RHEV 3.3.0 upgrade.
             Please read the following knowledge article for known
   issues and
             updated instructions before proceeding with the upgrade.
             RHEV 3.3.0 Upgrade Guide: Tips, Considerations and Roll-
   back Issues
                 https://access.redhat.com/site/articles/408623
             Would you like to continue with the upgrade? (Yes, No)
   [Yes]:
   ```

4. Remove Red Hat Enterprise Virtualization Manager 3.2 channels and entitlements to ensure the system does not use any Red Hat Enterprise Virtualization Manager 3.2 packages.

   ### Subscription Manager

   Use the **yum-config-manager** command to disable the Red Hat Enterprise Virtualization 3.2 repository in your **yum** configuration.

   ```
   # yum-config-manager --disable rhel-6-server-rhevm-3.2-rpms
   ```

   ### Red Hat Network Classic

Use the **rhn-channel** command or the Red Hat Network web interface to remove the **Red Hat Enterprise Virtualization Manager (v.3.2 x86_64)** channels.

```
# rhn-channel --remove --channel=rhel-x86_64-server-6-rhevm-3.2
```

5. Run the following command to ensure all packages related to Red Hat Enterprise Virtualization are up to date:

```
# yum update
```

In particular, if you are using the JBoss Application Server from JBoss Enterprise Application Platform 6.0.1, you must run the above command to upgrade to Enterprise Application Platform 6.1.

Result

Red Hat Enterprise Virtualization Manager has been upgraded. To take full advantage of all Red Hat Enterprise Virtualization 3.3 features you must also:

- Ensure all of your virtualization hosts are up to date and running the most recent Red Hat Enterprise Linux packages or Hypervisor images.
- Change all of your clusters to use compatibility version 3.3.
- Change all of your data centers to use compatibility version 3.3.

Report a bug

6.3. Upgrading to Red Hat Enterprise Virtualization Manager 3.2

6.3.1. Upgrading to Red Hat Enterprise Virtualization Manager 3.2

Prerequisites:

- Section 6.4.1, “Upgrading to Red Hat Enterprise Virtualization Manager 3.1”

Summary

Upgrading Red Hat Enterprise Virtualization Manager to version 3.2 is performed using the **rhevm-upgrade** command. Virtualization hosts, and the virtual machines running upon them, will continue to operate independently while the Manager is being upgraded. Once the Manager upgrade is complete you will be able to upgrade your hosts, if you haven't already, to the latest versions of Red Hat Enterprise Linux and Red Hat Enterprise Virtualization Hypervisor.

Important

Users of Red Hat Enterprise Virtualization 3.0 must migrate to Red Hat Enterprise Virtualization 3.1 before attempting this upgrade.

Note

In the event that the upgrade fails the **rhevm-upgrade** command will attempt to roll your Red Hat Enterprise Virtualization Manager installation back to its previous state. Where this also fails detailed instructions for manually restoring the installation are displayed.

Procedure 6.8. Upgrading to Red Hat Enterprise Virtualization Manager 3.2

1. **Add Red Hat Enterprise Virtualization 3.2 Subscription**

   Ensure that the system is subscribed to the required channels and entitlements to receive Red Hat Enterprise Virtualization Manager 3.2 packages. This procedure assumes that the system is already subscribed to required channels and entitlements to receive Red Hat Enterprise Virtualization 3.1 packages. These must also be available to complete the upgrade process.

   ### Certificate-based Red Hat Network

   The Red Hat Enterprise Virtualization 3.2 packages are provided by the **rhel-6-server-rhevm-3.2-rpms** repository associated with the **Red Hat Enterprise Virtualization** entitlement. Use the **yum-config-manager** command to enable the repository in your **yum** configuration. The **yum-config-manager** command must be run while logged in as the **root** user.

   ```
   # yum-config-manager --enable rhel-6-server-rhevm-3.2-rpms
   ```

   ### Red Hat Network Classic

   The Red Hat Enterprise Virtualization 3.2 packages are provided by the **Red Hat Enterprise Virtualization Manager (v.3.2 x86_64)** channel, also referred to as **rhel-x86_64-server-6-rhevm-3.2** in Red Hat Network Classic.

   ```
   rhn-channel --add --channel=rhel-x86_64-server-6-rhevm-3.2
   ```

   Use the **rhn-channel** command, or the Red Hat Network Web Interface, to subscribe to the **Red Hat Enterprise Virtualization Manager (v.3.2 x86_64)** channel.

2. **Remove Enterprise Virtualization 3.1 Subscription**

   Ensure that the system does not use any Red Hat Enterprise Virtualization Manager 3.1 packages by removing the Red Hat Enterprise Vitulization Manager 3.1 channels and entitlements.

   ### Certificate-based Red Hat Network

   Use the **yum-config-manager** command to disable the Red Hat Enterprise Virtualization 3.1 repository in your **yum** configuration. The **yum-config-manager** command must be run while logged in as the **root** user.

   ```
   # yum-config-manager --disablerepo=rhel-6-server-rhevm-3.1-rpms
   ```

   ### Red Hat Network Classic

Use the **rhn-channel** command, or the Red Hat Network Web Interface, to remove the **Red Hat Enterprise Virtualization Manager (v.3.1 x86_64)** channels.

```
# rhn-channel --remove --channel=rhel-6-server-rhevm-3.1
```

3. **Update the *rhevm-setup* Package**

   To ensure that you have the most recent version of the **rhevm-upgrade** command installed you must update the *rhevm-setup* package. Log in as the **root** user and use **yum** to update the *rhevm-setup* package.

   ```
   # yum update rhevm-setup
   ```

4. **Run the rhevm-upgrade Command**

   To upgrade Red Hat Enterprise Virtualization Manager run the **rhevm-upgrade** command. You must be logged in as the **root** user to run this command.

   ```
   # rhevm-upgrade
   Loaded plugins: product-id, rhnplugin
   Info: RHEV Manager 3.1 to 3.2 upgrade detected
   Checking pre-upgrade conditions...(This may take several minutes)
   ```

5. If the *ipa-server* package is installed then an error message is displayed. Red Hat Enterprise Virtualization Manager 3.2 does not support installation on the same machine as Identity Management (IdM).

   ```
   Error: IPA was found to be installed on this machine. Red Hat
   Enterprise Virtualization Manager 3.2 does not support installing
   IPA on the same machine. Please remove ipa packages before you
   continue.
   ```

   To resolve this issue you must migrate the IdM configuration to another system before re-attempting the upgrade. For further information see https://access.redhat.com/knowledge/articles/233143.

Result

Your Red Hat Enterprise Virtualization Manager installation has now been upgraded. To take full advantage of all Red Hat Enterprise Virtualization 3.2 features you must also:

- Ensure that all of your virtualization hosts are up to date and running the most recent Red Hat Enterprise Linux packages or Hypervisor images.
- Change all of your clusters to use compatibility version 3.2.
- Change all of your data centers to use compatibility version 3.2.

Report a bug

6.4. Upgrading to Red Hat Enterprise Virtualization Manager 3.1

6.4.1. Upgrading to Red Hat Enterprise Virtualization Manager 3.1

Summary

Upgrading Red Hat Enterprise Virtualization Manager to version 3.1 is performed using the **rhevm-upgrade** command. Virtualization hosts, and the virtual machines running upon them, will continue to operate independently while the Manager is being upgraded. Once the Manager upgrade is complete you will be able to upgrade your hosts, if you haven't already, to the latest versions of Red Hat Enterprise Linux and Red Hat Enterprise Virtualization Hypervisor.

Important

Refer to https://access.redhat.com/knowledge/articles/269333 for an up to date list of tips and considerations to be taken into account when upgrading to Red Hat Enterprise Virtualization 3.1.

Important

Users of Red Hat Enterprise Virtualization 2.2 must migrate to Red Hat Enterprise Virtualization 3.0 before attempting this upgrade. For information on migrating from Red Hat Enterprise Virtualization 2.2 to Red Hat Enterprise Virtualization 3.0, refer to https://access.redhat.com/knowledge/techbriefs/migrating-red-hat-enterprise-virtualization-manager-version-22-30.

Note

In the event that the upgrade fails the **rhevm-upgrade** command will attempt to roll your Red Hat Enterprise Virtualization Manager installation back to its previous state. Where this also fails detailed instructions for manually restoring the installation are displayed.

Procedure 6.9. Upgrading to Red Hat Enterprise Virtualization Manager 3.1

1. **Red Hat JBoss Enterprise Application Platform 6 Subscription**

   Ensure that the system is subscribed to the required channels and entitlements to receive Red Hat JBoss Enterprise Application Platform 6 packages. Red Hat JBoss Enterprise Application Platform 6 is a required dependency of Red Hat Enterprise Virtualization Manager 3.1.

   #### Certificate-based Red Hat Network

   The Red Hat JBoss Enterprise Application Platform 6 packages are provided by the **Red Hat JBoss Enterprise Application Platform** entitlement in certificate-based Red Hat Network.

   Use the **subscription-manager** command to ensure that the system is subscribed to the **Red Hat JBoss Enterprise Application Platform** entitlement.

   ```
   # subscription-manager list
   ```

   #### Red Hat Network Classic

The Red Hat JBoss Enterprise Application Platform 6 packages are provided by the **Red Hat JBoss Application Platform (v 6) for 6Server x86_64** channel, also referred to as **jbappplatform-6-x86_64-server-6-rpm**, in Red Hat Network Classic. The Channel Entitlement Name for this channel is **Red Hat JBoss Enterprise Application Platform (v 4, zip format)**.

Use the **rhn-channel** command, or the Red Hat Network Web Interface, to subscribe to the **Red Hat JBoss Application Platform (v 6) for 6Server x86_64** channel.

2. **Add Red Hat Enterprise Virtualization 3.1 Subscription**

   Ensure that the system is subscribed to the required channels and entitlements to receive Red Hat Enterprise Virtualization Manager 3.1 packages.

   ### Certificate-based Red Hat Network

   The Red Hat Enterprise Virtualization 3.1 packages are provided by the **rhel-6-server-rhevm-3.1-rpms** repository associated with the **Red Hat Enterprise Virtualization** entitlement. Use the **yum-config-manager** command to enable the repository in your **yum** configuration. The **yum-config-manager** command must be run while logged in as the **root** user.

   ```
   # yum-config-manager --enable rhel-6-server-rhevm-3.1-rpms
   ```

   ### Red Hat Network Classic

   The Red Hat Enterprise Virtualization 3.1 packages are provided by the **Red Hat Enterprise Virtualization Manager (v.3.1 x86_64)** channel, also referred to as **rhel-x86_64-server-6-rhevm-3.1** in Red Hat Network Classic.

   Use the **rhn-channel** command, or the Red Hat Network Web Interface, to subscribe to the **Red Hat Enterprise Virtualization Manager (v.3.1 x86_64)** channel.

3. **Remove Red Hat Enterprise Virtualization 3.0 Subscription**

   Ensure that the system does not use any Red Hat Enterprise Virtualization Manager 3.0 packages by removing the Red Hat Enterprise Virtualization Manager 3.0 channels and entitlements.

   ### Certificate-based Red Hat Network

   Use the **yum-config-manager** command to disable the Red Hat Enterprise Virtualization 3.0 repositories in your **yum** configuration. The **yum-config-manager** command must be run while logged in as the **root** user.

   ```
   # yum-config-manager --disablerepo=rhel-6-server-rhevm-3-rpms
   ```

   ```
   # yum-config-manager --disablerepo=jb-eap-5-for-rhel-6-server-rpms
   ```

   ### Red Hat Network Classic

   Use the **rhn-channel** command, or the Red Hat Network Web Interface, to remove the **Red Hat Enterprise Virtualization Manager (v.3.0 x86_64)** channels.

   ```
   # rhn-channel --remove --channel=rhel-6-server-rhevm-3
   ```

```
# rhn-channel --remove --channel=jbappplatform-5-x86_64-server-6-
rpm
```

4. **Update the *rhevm-setup* Package**

   To ensure that you have the most recent version of the **rhevm-upgrade** command installed you must update the *rhevm-setup* package. Log in as the **root** user and use **yum** to update the *rhevm-setup* package.

   ```
   # yum update rhevm-setup
   ```

5. **Run the rhevm-upgrade Command**

   To upgrade Red Hat Enterprise Virtualization Manager run the **rhevm-upgrade** command. You must be logged in as the **root** user to run this command.

   ```
   # rhevm-upgrade
   Loaded plugins: product-id, rhnplugin
   Info: RHEV Manager 3.0 to 3.1 upgrade detected
   Checking pre-upgrade conditions...(This may take several minutes)
   ```

6. If the *ipa-server* package is installed then an error message is displayed. Red Hat Enterprise Virtualization Manager 3.1 does not support installation on the same machine as Identity Management (IdM).

   ```
   Error: IPA was found to be installed on this machine. Red Hat
   Enterprise Virtualization Manager 3.1 does not support installing
   IPA on the same machine. Please remove ipa packages before you
   continue.
   ```

   To resolve this issue you must migrate the IdM configuration to another system before re-attempting the upgrade. For further information see https://access.redhat.com/knowledge/articles/233143.

7. A list of packages that depend on Red Hat JBoss Enterprise Application Platform 5 is displayed. These packages must be removed to install Red Hat JBoss Enterprise Application Platform 6, required by Red Hat Enterprise Virtualization Manager 3.1.

   ```
    Warning: the following packages will be removed if you proceed
   with the upgrade:

       * objectweb-asm

    Would you like to proceed? (yes|no):
   ```

   You must enter **yes** to proceed with the upgrade, removing the listed packages.

Result

Your Red Hat Enterprise Virtualization Manager installation has now been upgraded. To take full advantage of all Red Hat Enterprise Virtualization 3.1 features you must also:

- Ensure that all of your virtualization hosts are up to date and running the most recent Red Hat Enterprise Linux packages or Hypervisor images.

- Change all of your clusters to use compatibility version 3.1.
- Change all of your data centers to use compatibility version 3.1.

See Also:

- Section 6.1.5, “Updating Red Hat Enterprise Virtualization Hypervisors”
- Section 6.1.6, “Updating Red Hat Enterprise Linux Virtualization Hosts”
- Section 6.5.2, “Changing the Cluster Compatibility Version”
- Section 6.5.3, “Changing the Data Center Compatibility Version”

Report a bug

6.5. Post-upgrade Tasks

6.5.1. Features Requiring a Compatibility Upgrade to Red Hat Enterprise Virtualization 3.3

Some of the new features in Red Hat Enterprise Virtualization are only available if your data centers, clusters, and storage have a compatibility version of **3.3**.

Table 6.1. Features Requiring a Compatibility Upgrade to Red Hat Enterprise Virtualization 3.3

Feature	Description
Libvirt-to-libvirt virtual machine migration	Perform virtual machine migration using libvirt-to-libvirt communication. This is safer, more secure, and has less host configuration requirements than native KVM migration, but has a higher overhead on the host CPU.
Isolated network to carry virtual machine migration traffic	Separates virtual machine migration traffic from other traffic types, like management and display traffic. Reduces chances of migrations causing a network flood that disrupts other important traffic types.
Define a gateway per logical network	Each logical network can have a gateway defined as separate from the management network gateway. This allows more customizable network topologies.
Snapshots including RAM	Snapshots now include the state of a virtual machine's memory as well as disk.
Optimized iSCSI device driver for virtual machines	Virtual machines can now consume iSCSI storage as virtual hard disks using an optimized device driver.
Host support for MOM management of memory overcommitment	MOM is a policy-driven tool that can be used to manage overcommitment on hosts. Currently MOM supports control of memory ballooning and KSM.
GlusterFS data domains.	Native support for the GlusterFS protocol was added as a way to create storage domains, allowing Gluster data centers to be created.

Feature	Description
Custom device property support	In addition to defining custom properties of virtual machines, you can also define custom properties of virtual machine devices.
Multiple monitors using a single virtual PCI device	Drive multiple monitors using a single virtual PCI device, rather than one PCI device per monitor.
Updatable storage server connections	It is now possible to edit the storage server connection details of a storage domain.
Check virtual hard disk alignment	Check if a virtual disk, the filesystem installed on it, and its underlying storage are aligned. If it is not aligned, there may be a performance penalty.
Extendable virtual machine disk images	You can now grow your virtual machine disk image when it fills up.
OpenStack Image Service integration	Red Hat Enterprise Virtualization supports the OpenStack Image Service. You can import images from and export images to an Image Service repository.
Gluster hook support	You can manage Gluster hooks, which extend volume life cycle events, from Red Hat Enterprise Virtualization Manager.
Gluster host UUID support	This feature allows a Gluster host to be identified by the Gluster server UUID generated by Gluster in addition to identifying a Gluster host by IP address.
Network quality of service (QoS) support	Limit the inbound and outbound network traffic at the virtual NIC level.
Cloud-Init support	Cloud-Init allows you to automate early configuration tasks in your virtual machines, including setting hostnames, authorized keys, and more.

Report a bug

6.5.2. Changing the Cluster Compatibility Version

Prerequisites:

- Section 6.1.5, “Updating Red Hat Enterprise Virtualization Hypervisors”
- Section 6.1.6, “Updating Red Hat Enterprise Linux Virtualization Hosts”

Summary

Red Hat Enterprise Virtualization clusters have a compatibility version. The cluster compatibility version indicates the features of Red Hat Enterprise Virtualization supported by all of the hosts in the cluster. The cluster compatibility is set according to the version of the least capable host operating system in the cluster.

Note

To change the cluster compatibility version, you must have first updated all the hosts in your cluster to a level that supports your desired compatibility level.

Procedure 6.10. Changing the Cluster Compatibility Version

1. Log in to the Administration Portal as the administrative user. By default this is the **admin** user.
2. Click the **Clusters** tab.
3. Select the cluster that you wish to change from the list displayed. If the list of clusters is too long to filter visually then perform a search to locate the desired cluster.
4. Click the **Edit** button.
5. Change the **Compatibility Version** to the desired value.
6. Click **OK**.

Result

You have updated the compatibility version of the cluster. Once you have updated the compatibility version of all clusters in a data center, then you are also able to change the compatibility version of the data center itself.

Report a bug

6.5.3. Changing the Data Center Compatibility Version

Prerequisites:

- Section 6.5.2, “Changing the Cluster Compatibility Version”

Summary

Red Hat Enterprise Virtualization data centers have a compatibility version. The compatibility version indicates the version of Red Hat Enterprise Virtualization that the data center is intended to be compatible with. All clusters in the data center must support the desired compatibility level.

Procedure 6.11. Changing the Data Center Compatibility Version

1. Log in to the Administration Portal as the administrative user. By default this is the **admin** user.
2. Click the **Data Centers** tab.
3. Select the data center that you wish to change from the list displayed. If the list of data centers is too long to filter visually then perform a search to locate the desired data center.
4. Click the **Edit** button.
5. Change the **Compatibility Version** to the desired value.
6. Click **OK**.

Result

You have updated the compatibility version of the data center.

Report a bug

Part III. Installing Virtualization Hosts

Chapter 7. Introduction to Virtualization Hosts

7.1. Workflow Progress — Installing Virtualization Hosts

Report a bug

7.2. Introduction to Virtualization Hosts

Red Hat Enterprise Virtualization supports both virtualization hosts which run the Red Hat Enterprise Virtualization Hypervisor, and those which run Red Hat Enterprise Linux. Both types of virtualization host are able to coexist in the same Red Hat Enterprise Virtualization environment.

Prior to installing virtualization hosts you should ensure that:

- all virtualization hosts meet the hardware requirements, and
- you have successfully completed installation of the Red Hat Enterprise Virtualization Manager.

Additionally you may have chosen to install the Red Hat Enterprise Virtualization Manager Reports. This is not mandatory and is not required to commence installing virtualization hosts. Once you have completed the above tasks you are ready to install virtualization hosts.

Important

It is recommended that you install at least two virtualization hosts and attach them to the Red Hat Enterprise Virtualization environment. Where you attach only one virtualization host you will be unable to access features such as migration which require redundant hosts.

Important

The Red Hat Enterprise Virtualization Hypervisor is a closed system. Use a Red Hat Enterprise Linux host if additional rpms are required for your environment.

- Chapter 8, *Installing Red Hat Enterprise Virtualization Hypervisor Hosts*
- Chapter 9, *Installing Red Hat Enterprise Linux Hosts*

See Also:

- Section 2.2.3, “Virtualization Host Hardware Requirements”

Report a bug

Chapter 8. Installing Red Hat Enterprise Virtualization Hypervisor Hosts

8.1. Red Hat Enterprise Virtualization Hypervisor Installation Overview

Before commencing Hypervisor installation you must be aware that:

- The Red Hat Enterprise Virtualization Hypervisor *must* be installed on a physical server. It must not be installed in a Virtual Machine.
- The installation process will reconfigure the selected storage device and destroy all data. Therefore, ensure that any data to be retained is successfully backed up before proceeding.
- All Hypervisors in an environment must have unique hostnames and IP addresses, in order to avoid network conflicts.
- Instructions for using Network (PXE) Boot to install the Hypervisor are contained in the *Red Hat Enterprise Linux - Installation Guide*, available at https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux.
- Red Hat Enterprise Virtualization Hypervisors can use Storage Attached Networks (SANs) and other network storage for storing virtualized guest images. However, a local storage device is required for installing and booting the Hypervisor.

Note

Red Hat Enterprise Virtualization Hypervisor installations can be automated or conducted without interaction. This type of installation is only recommended for advanced users.

Report a bug

8.2. Installing the Red Hat Enterprise Virtualization Hypervisor Packages

Summary

The **`Red Hat Enterprise Virtualization Hypervisor (v.6 x86_64)`** Red Hat Network channel contains the Hypervisor packages. The Hypervisor itself is contained in the *rhev-hypervisor6* package. Additional tools supporting USB and PXE installations are also installed as a dependency. You must install the Hypervisor packages on the system that you intend to use to create Hypervisor boot media.

Procedure 8.1. Installing the Red Hat Enterprise Virtualization Hypervisor Packages

1. A. **Subscribing to Download the Hypervisor using Certificate-Based RHN**

    a. **Identify Available Entitlement Pools**

    To subscribe the system to Red Hat Enterprise Virtualization, you must locate the identifier for the relevant entitlement pool. Use the **`list`** action of the **`subscription-manager`** to find these:

To identify available subscription pools for **Red Hat Enterprise Virtualization**, use the command:

```
# subscription-manager list --available | grep -A8 "Red Hat
Enterprise Virtualization"
```

b. **Attach Entitlement Pools to the System**

Using the pool identifiers located in the previous step, attach the **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** entitlements to the system. Use the ***attach*** parameter of the **subscription-manager** command, replacing *[POOLID]* with each of the pool identifiers:

```
# subscription-manager attach --pool=[POOLID]
```

B. **Subscribing to Download the Hypervisor using RHN Classic**

a. Log on to Red Hat Network (http://rhn.redhat.com).

b. Move the mouse cursor over the **Subscriptions** link at the top of the page, and then click **Registered Systems** in the menu that appears.

c. Select the system to which you are adding channels from the list presented on the screen, by clicking the name of the system.

d. Click **Alter Channel Subscriptions** in the **Subscribed Channels** section of the screen.

e. Select the **Red Hat Enterprise Virtualization Hypervisor (v.6 x86_64)** channel from the list presented on the screen, then click the **Change Subscription** button to finalize the change.

2. Log in to the system on which the Red Hat Enterprise Virtualization Manager is installed. You must log in as the **root** user.

3. Use **yum** to install *rhev-hypervisor6*:

```
# yum install rhev-hypervisor6
```

4. Use **yum** to install *livecd-tools*:

```
# yum install livecd-tools
```

Result

The Hypervisor ISO image is installed into the **/usr/share/rhev-hypervisor/** directory. The **livecd-iso-to-disk** and **livecd-iso-to-pxeboot** scripts are installed to the **/usr/bin** directory.

Note

Red Hat Enterprise Linux 6.2 and later allows more than one version of the ISO image to be installed at one time. As such, **/usr/share/rhev-hypervisor/rhev-hypervisor.iso** is now a symbolic link to a uniquely-named version of the Hypervisor ISO image, such as **/usr/share/rhev-hypervisor/rhev-hypervisor-6.4-20130321.0.el6ev.iso**. Different versions of the image can now be installed alongside each other, allowing administrators to run and maintain a cluster on a previous version of the Hypervisor while upgrading another cluster for testing.

Additionally, the symbolic link **/usr/share/rhev-hypervisor/rhevh-latest-6.iso**, is created. This links also targets the most recently installed version of the Red Hat Enterprise Virtualization ISO image.

Report a bug

8.3. Preparing Hypervisor Installation Media

8.3.1. Preparing USB Hypervisor Installation Media

8.3.1.1. Preparing a Hypervisor USB Storage Device

Red Hat Enterprise Virtualization Hypervisor is able to install itself onto USB storage devices or solid state disks. However, the initial boot/install USB device must be a separate device from the installation target. Network booting with PXE and tftp provides the greatest flexibility and scalability. For environments where network restrictions prevent network booting, or for systems without PXE capable network interface cards, a local media installation such as CD-ROM or USB is necessary. Booting from USB storage devices is a useful alternative to booting from CD, for systems without CD-ROM drives.

Note

Not all systems support booting from a USB storage device. Ensure that your system's BIOS supports booting from USB storage devices before proceeding.

See Also:

- Section 8.3.1.2, “Preparing USB Installation Media Using livecd-iso-to-disk”
- Section 8.3.1.3, “Preparing USB Installation Media Using dd”

Report a bug

8.3.1.2. Preparing USB Installation Media Using livecd-iso-to-disk

Summary

The **livecd-iso-to-disk** utility included in the *livecd-tools* package can be used to write a Hypervisor or other disk image to a USB storage device. Once a Hypervisor disk image has been written to a USB storage device with this utility, systems that support booting via USB can boot the Hypervisor using the USB storage device.

The basic syntax for the **livecd-iso-to-disk** utility is as follows:

```
# livecd-iso-to-disk [image] [device]
```

The *[device]* parameter is the path to the USB storage device on which to write the disk image. The *[image]* parameter is the path and file name of the disk image. The default Hypervisor image location is **/usr/share/rhev-hypervisor/rhev-hypervisor.iso**. The **livecd-iso-to-disk** utility requires devices to be formatted with the **FAT** or **EXT3** file system.

Note

The **livecd-iso-to-disk** utility uses a **FAT** or **EXT3** formatted partition or block device.

Note

USB storage devices are sometimes formatted without a partition table. In this case, use a generic identifier for the storage device such as **/dev/sdb**. When a USB storage device is formatted with a partition table, use the path name to the device, such as **/dev/sdb1**.

Procedure 8.2. Preparing USB Installation Media Using livecd-iso-to-disk

1. Install the *rhev-hypervisor* package to download the latest version of the Hypervisor.
2. Use the **livecd-iso-to-disk** utility to copy the disk image, located in the **/usr/share/rhev-hypervisor/** directory, to the USB storage device. The ***--format*** parameter formats the USB device. The ***--reset-mbr*** parameter initializes the Master Boot Record (MBR).

Example 8.1. Use of livecd-iso-to-disk

This example demonstrates the use of **livecd-iso-to-disk** to write to a USB storage device named **/dev/sdc** and make the USB storage device bootable.

```
# livecd-iso-to-disk --format --reset-mbr /usr/share/rhev-
hypervisor/rhev-hypervisor.iso /dev/sdc
Verifying image...
/usr/share/rhev-hypervisor/rhev-hypervisor.iso:
eccc12a0530b9f22e5ba62b848922309
Fragment sums:
8688f5473e9c176a73f7a37499358557e6c397c9ce2dafb5eca5498fb586
Fragment count: 20
Press [Esc] to abort check.
Checking: 100.0%

The media check is complete, the result is: PASS.
```

```
It is OK to use this media.

WARNING: THIS WILL DESTROY ANY DATA ON /dev/sdc!!!
Press Enter to continue or ctrl-c to abort

/dev/sdc: 2 bytes were erased at offset 0x000001fe (dos): 55 aa
Waiting for devices to settle...
mke2fs 1.42.7 (21-Jan-2013)
Filesystem label=LIVE
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
488640 inodes, 1953280 blocks
97664 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2000683008
60 block groups
32768 blocks per group, 32768 fragments per group
8144 inodes per group
Superblock backups stored on blocks:
        32768, 98304, 163840, 229376, 294912, 819200, 884736,
1605632

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

Copying live image to target device.
squashfs.img
   163360768 100%  184.33MB/s    0:00:00 (xfer#1, to-check=0/1)

sent 163380785 bytes  received 31 bytes  108920544.00 bytes/sec
total size is 163360768  speedup is 1.00
osmin.img
        4096 100%    0.00kB/s    0:00:00 (xfer#1, to-check=0/1)

sent 4169 bytes  received 31 bytes  8400.00 bytes/sec
total size is 4096  speedup is 0.98
Updating boot config file
Installing boot loader
/media/tgttmp.q6aZdS/syslinux is device /dev/sdc
Target device is now set up with a Live image!
```

Result

The Hypervisor disk image has been written to the USB storage device. You can now use it to boot a system and install the Hypervisor.

Report a bug

8.3.1.3. Preparing USB Installation Media Using dd

The **dd** command can also be used to install a hypervisor onto a USB storage device. Media created with the command can boot the Hypervisor on systems which support booting via USB. Red Hat Enterprise Linux provides **dd** as part of the *coreutils* package. Versions of **dd** are also available on a wide variety of Linux and Unix operating systems.

Windows users are able to obtain the **dd** command through installation of **Red Hat Cygwin**, a free Linux-like environment for Windows.

The basic **dd** command usage follows this structure:

```
# dd if=image of=device
```

Where the ***device*** parameter is the device name of the USB storage device to install to. The ***image*** parameter is a ISO image of the Hypervisor. The default hypervisor image location is **/usr/share/rhev-hypervisor/rhev-hypervisor.iso**. The **dd** command does not make assumptions as to the format of the device as it performs a low-level copy of the raw data in the selected image.

See Also:

- Section 8.3.1.4, “Preparing USB Installation Media Using dd on Linux Systems”
- Section 8.3.1.5, “Preparing USB Installation Media Using dd on Windows Systems”

Report a bug

8.3.1.4. Preparing USB Installation Media Using dd on Linux Systems

Summary

The **dd** command available on most Linux systems is suitable for creating USB installation media, to boot and install the Hypervisor.

Procedure 8.3. Preparing USB Installation Media using dd on Linux Systems

1. Install the *rhev-hypervisor* package.

   ```
   # yum install rhev-hypervisor
   ```

2. Use the **dd** command to copy the image file to the disk.

   Example 8.2. Use of dd

   This example uses a USB storage device named **/dev/sdc**.

   ```
   # dd if=/usr/share/rhev-hypervisor/rhev-hypervisor.iso
   of=/dev/sdc
   243712+0 records in
   243712+0 records out
   124780544 bytes (125 MB) copied, 56.3009 s, 2.2 MB/s
   ```

Warning

The **dd** command will overwrite all data on the device specified for the ***of*** parameter. Any existing data on the device will be destroyed. Ensure that the correct device is specified and that it contains no valuable data before invocation of the **dd** command.

Result

The USB storage device is ready for use as Hypervisor installation media.

Report a bug

8.3.1.5. Preparing USB Installation Media Using dd on Windows Systems

Summary

The **dd** command, available on Windows systems with Red Hat Cygwin installed, is suitable for creating USB installation media to boot and install the Hypervisor.

Procedure 8.4. Preparing USB Installation Media using dd on Windows Systems

1. Access http://www.redhat.com/services/custom/cygwin/ and click the **Red Hat Cygwin official installation utility** link. The **rhsetup.exe** executable will download.
2. As the **Administrator** user run the downloaded **rhsetup.exe** executable. The **Red Hat Cygwin** installer will display.
3. Follow the prompts to complete a standard installation of **Red Hat Cygwin**. The *Coreutils* package within the *Base* package group provides the **dd** utility. This is automatically selected for installation.
4. Copy the **rhev-hypervisor.iso** file downloaded from **Red Hat Network** to **C:\rhev-hypervisor.iso**.
5. As the **Administrator** user run **Red Hat Cygwin** from the desktop. A terminal window will appear.

Important

On the **Windows 7** and **Windows Server 2008** platforms it is necessary to right click the **Red Hat Cygwin** icon and select the **Run as Administrator...** option to ensure the application runs with the correct permissions.

6. In the terminal run **cat /proc/partitions** to see the drives and partitions currently visible to the system.

Example 8.3. View of Disk Partitions Attached to System

```
Administrator@test /
$ cat /proc/partitions
major minor  #blocks  name
    8     0  15728640 sda
```

```
    8     1    102400 sda1
    8     2  15624192 sda2
```

7. Plug the USB storage device which is to be used as the media for the Hypervisor installation into the system. Re-run the **cat /proc/partitions** command and compare the output to that of the previous run. A new entry will appear which designates the USB storage device.

Example 8.4. View of Disk Partitions Attached to System

```
Administrator@test /
$ cat /proc/partitions
major minor  #blocks  name
    8     0  15728640 sda
    8     1    102400 sda1
    8     2  15624192 sda2
    8    16    524288 sdb
```

8. Use the **dd** command to copy the **rhev-hypervisor.iso** file to the disk. The example uses a USB storage device named **/dev/*sdb***. Replace *sdb* with the correct device name for the USB storage device to be used.

Example 8.5. Use of dd Command Under Red Hat Cygwin

```
Administrator@test /
$ dd if=/cygdrive/c/rhev-hypervisor.iso of=/dev/sdb& pid=$!
```

The provided command starts the transfer in the background and saves the process identifier so that it can be used to monitor the progress of the transfer. Refer to the next step for the command used to check the progress of the transfer.

Warning

The **dd** command will overwrite all data on the device specified for the ***of*** parameter. Any existing data on the device will be destroyed. Ensure that the correct device is specified and that it contains no valuable data before invocation of the **dd** command.

9. Transfer of the ISO file to the USB storage device with the version of **dd** included with **Red Hat Cygwin** can take significantly longer than the equivalent on other platforms.

   To check the progress of the transfer in the same terminal window that the process was started in send it the **USR1** signal. This can be achieved by issuing the **kill** command in the terminal window as follows:

```
kill -USR1 $pid
```

10. When the transfer operation completes the final record counts will be displayed.

Example 8.6. Result of dd Initiated Copy

```
210944+0 records in
210944+0 records out
108003328 bytes (108 MB) copied, 2035.82 s, 53.1 kB/s

[1]+    Done            dd if=/cygdrive/c/rhev-hypervisor.iso
of=/dev/sdb
```

Result

The USB storage device is ready for use as Hypervisor installation media.

Report a bug

8.3.2. Preparing Optical Hypervisor Installation Media

Summary

Burn the Hypervisor image to a CD-ROM with the **wodim** command. The **wodim** command is part of the *wodim* package.

Procedure 8.5. Preparing Optical Hypervisor Installation Media

1. Verify that the *wodim* package is installed on the system.

   Example 8.7. Verify Installation of *wodim* Package

   ```
   # rpm -q wodim
   wodim-1.1.9-11.el6.x86_64
   ```

   If the package version is in the output the package is available.

   If nothing is listed, install *wodim*:

   ```
   # yum install wodim
   ```

2. Insert a blank CD-ROM or DVD into your CD or DVD writer.

3. Record the ISO file to the disc. The *wodim* command uses the following:

   ```
   wodim dev=device image
   ```

   This example uses the first CD-RW (**/dev/cdrw**) device available and the default hypervisor image location, **/usr/share/rhev-hypervisor/rhev-hypervisor.iso**.

   Example 8.8. Use of wodim Command

   ```
   # wodim dev=/dev/cdrw /usr/share/rhev-hypervisor/rhev-
   hypervisor.iso
   ```

Result

If no errors occurred, the Hypervisor is ready to boot. Errors sometimes occur during the recording process due to errors on the media itself. If this occurs insert another writable disk and repeat the command above.

The Hypervisor uses a program (**`isomd5sum`**) to verify the integrity of the installation media every time the Hypervisor is booted. If media errors are reported in the boot sequence you have a bad CD-ROM. Follow the procedure above to create a new CD-ROM or DVD.

Report a bug

8.3.3. Booting from Hypervisor Installation Media

8.3.3.1. Booting the Hypervisor from USB Installation Media

Summary

Booting a hypervisor from a USB storage device is similar to booting other live USB operating systems. Follow this procedure to boot the Hypervisor using USB installation media.

Procedure 8.6. Booting the Hypervisor from USB Installation Media

1. Enter the system's BIOS menu to enable USB storage device booting if not already enabled.
    a. Enable USB booting if this feature is disabled.
    b. Set booting USB storage devices to be first boot device.
    c. Shut down the system.
2. Insert the USB storage device that contains the Hypervisor boot image.
3. Restart the system.

Result

The Hypervisor boot process commences automatically.

See Also:

- Section 8.3.3.3, “Troubleshooting BIOS Settings and Boot Process”
- Section 8.3.3.4, “Choosing Hypervisor Boot Options”

Report a bug

8.3.3.2. Booting the Hypervisor from Optical Installation Media

Summary

Booting the Hypervisor from optical installation media requires the system to have a correctly defined BIOS boot configuration.

Procedure 8.7. Booting the Hypervisor from Optical Installation Media

1. Ensure that the system's BIOS is configured to boot from the CD-ROM or DVD-ROM drive first. For many systems this the default.

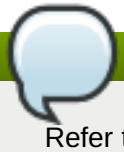

Note

Refer to your manufacturer's manuals for further information on modifying the system's BIOS boot configuration.

2. Insert the Hypervisor CD-ROM in the CD-ROM or DVD-ROM drive.
3. Reboot the system.

Result

The Hypervisor boot screen will be displayed.

See Also:

- Section 8.3.3.3, “Troubleshooting BIOS Settings and Boot Process”
- Section 8.3.3.4, “Choosing Hypervisor Boot Options”

Report a bug

8.3.3.3. Troubleshooting BIOS Settings and Boot Process

Summary

Before installing Red Hat Enterprise Virtualization Hypervisors it is necessary to verify the BIOS is correctly configured for the chosen installation method. Many motherboard and PC manufacturers disable different booting methods in the BIOS. Most BIOS chips boot from the following devices in order:

1. 3.5 inch diskette
2. CD-ROM or DVD device
3. Local hard disk

Many BIOS chips have disabled one or more of the following boot methods: USB storage devices, CD-ROMs, DVDs or network boot. To boot from your chosen method, enable the method or device and set that device as the first boot device in BIOS.

Most but not all motherboards support the boot methods described in this chapter. Consult the documentation for your motherboard or system to determine whether it is possible to use a particular boot method.

Warning

BIOS settings vary between manufacturers. Any specific examples of BIOS settings may be inaccurate for some systems. Due to this inconsistency, it is necessary to review the motherboard or system manufacturer's documentation.

Procedure 8.8. Troubleshooting BIOS Settings and Boot Process

Verify that your system is capable of running the Red Hat Enterprise Virtualization Hypervisor. Hypervisors require that virtualization extensions are present and enabled in the BIOS before installation proceeds.

1. Boot the Hypervisor from removable media. For example, a USB stick or CD-ROM.
2. When the message **`Automatic boot in 30 seconds...`** is displayed, and begins counting down from thirty, press any key to skip the automatic boot process.
3. Ensure the **`Install or Upgrade`** option is selected and press **`Tab`** to edit the boot parameters.
4. Add the **`rescue`** parameter to the list of boot parameters shown on the screen, then press **`Enter`**. This action will boot the Hypervisor in rescue mode.
5. Once the Hypervisor boots, verify your CPU contains the virtualization extensions with the following command:

   ```
   # grep -E "svm|vmx" /proc/cpuinfo
   ```

   Output displays if the processor has the hardware virtualization extensions.
6. Verify that the KVM modules load by default:

   ```
   # lsmod | grep kvm
   ```

Result

If the output includes **`kvm_intel`** or **`kvm_amd`** then the kvm hardware virtualization modules are loaded and the system meets the requirements. If the output does not include the required modules then you must check that your hardware supports the virtualization extensions and that they are enabled in the system's BIOS.

Report a bug

8.3.3.4. Choosing Hypervisor Boot Options

Summary

The Hypervisor boot menu presents a number of pre-configured boot modes, as well as the ability to add and remove additional boot options.

Procedure 8.9. Choosing Hypervisor Boot Options

1. Insert the Red Hat Enterprise Virtualization Hypervisor installation media.
2. Power on the system and ensure the system boots from the installation media.
3. The boot splash screen appears. If no input is provided, the Hypervisor installation will commence in 30 seconds, using default kernel parameters.

Figure 8.1. Boot Splash Screen

4. To modify the boot options, press any key. The boot menu will display.

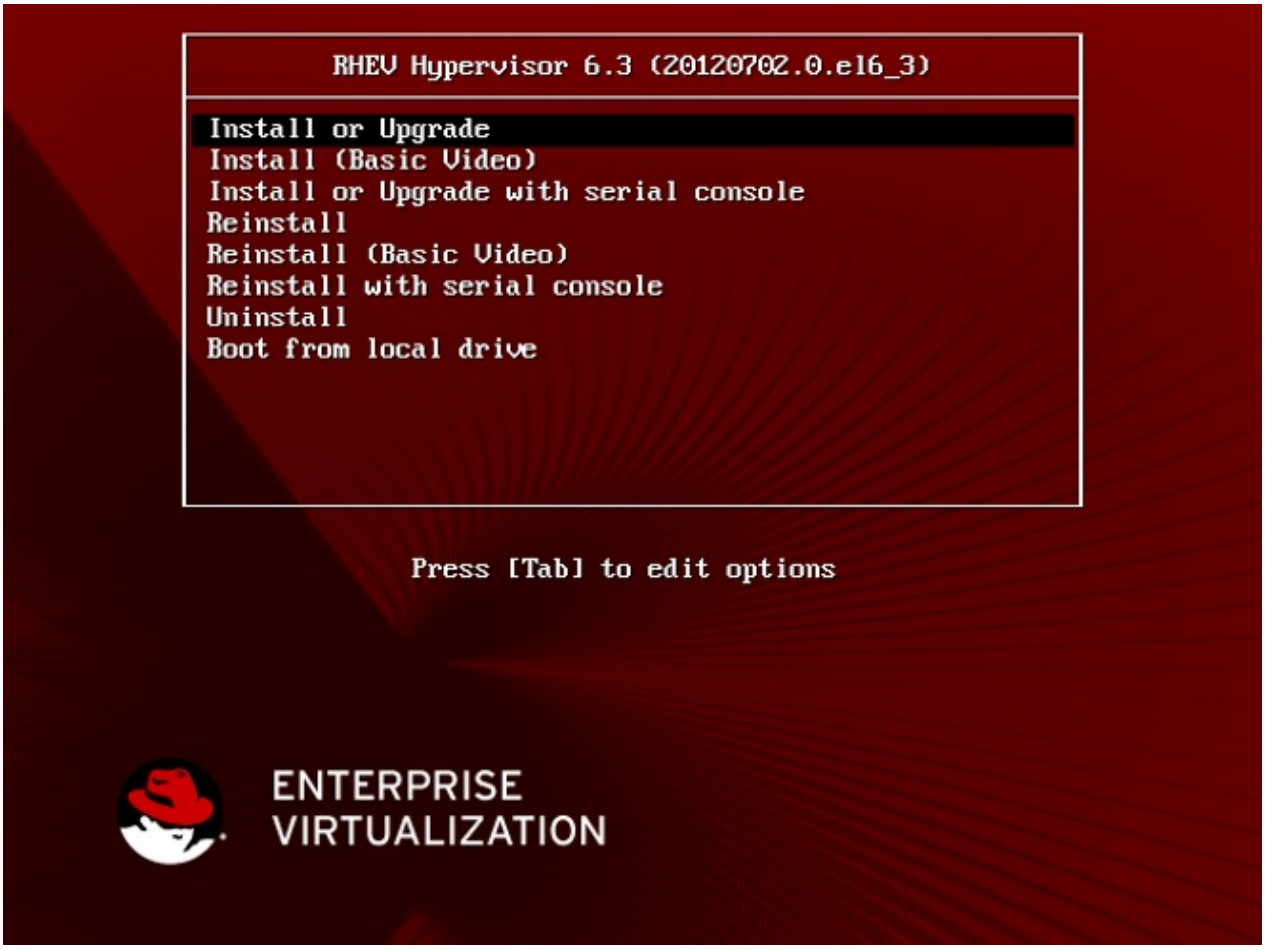

Figure 8.2. Boot Menu Screen

The following boot options are available:

Install or Upgrade

Boot the Hypervisor installer.

Install (Basic Video)

Install or Upgrade the Hypervisor, using basic video mode.

Install or Upgrade with Serial Console

Install or Upgrade the Hypervisor, with the console redirected to a serial device attached to **/dev/ttyS0**.

Reinstall

Reinstall the Hypervisor.

Reinstall (Basic Video)

Reinstall the Hypervisor, using basic video mode.

Reinstall with Serial Console

Reinstall the Hypervisor, with the console redirected to a serial device attached to **/dev/ttyS0**.

Boot from Local Drive

Boot the operating system installed on the first local drive.

Select the appropriate boot option from the boot menu.

5. A. Press the **Enter** key to boot the Hypervisor with the default kernel parameters for the option selected; or

   B. press the **Tab** key to edit the kernel parameters. In edit mode you are able to add or remove kernel parameters. Kernel parameters must be separated from each other by a space. Once the desired kernel parameters have been set press **Enter** to boot the system. Alternatively pressing **Esc** reverts any changes that you have made to the kernel parameters.

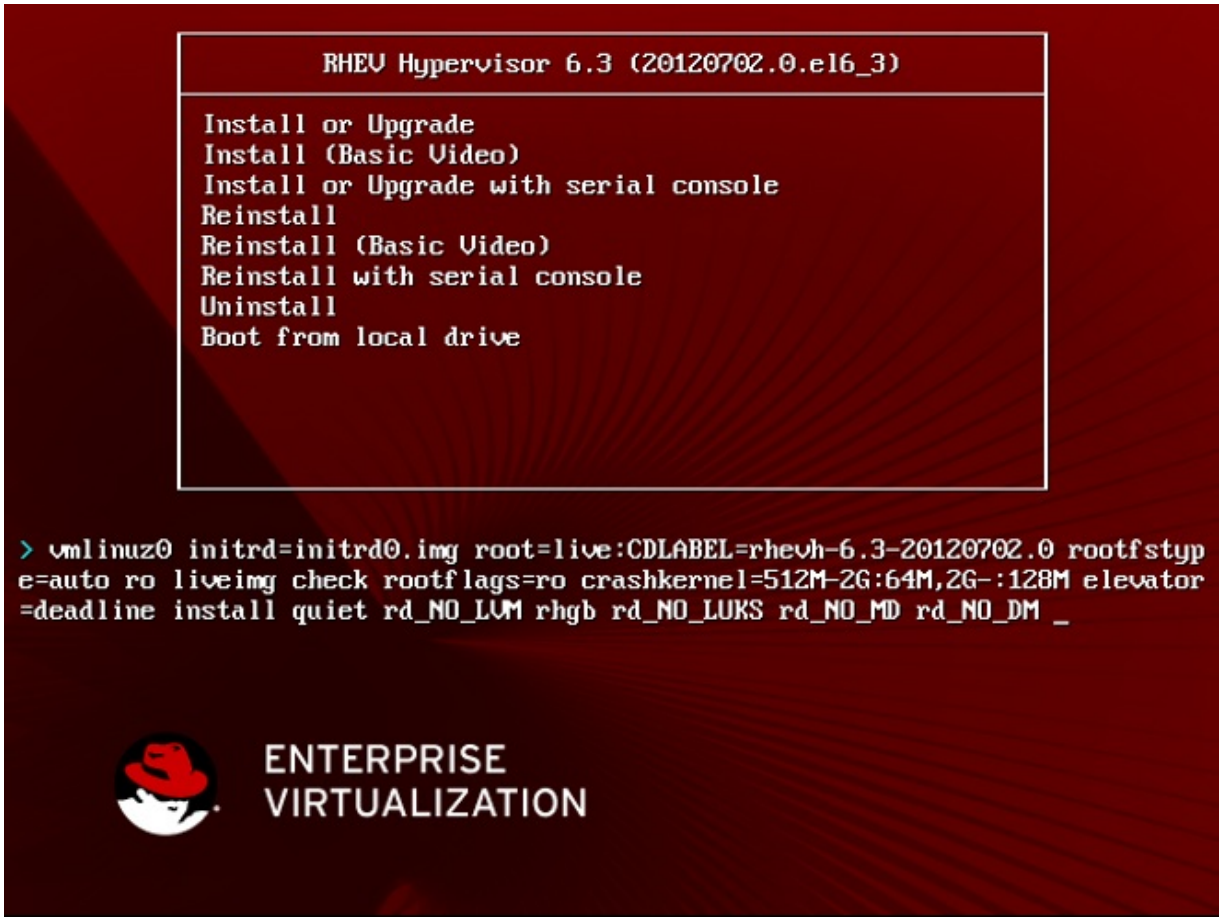

Figure 8.3. Boot Parameter Screen

Result

The Hypervisor boots with the provided boot options.

Report a bug

8.4. Installing the Hypervisor

8.4.1. Hypervisor Menu Actions

- The directional keys (**Up**, **Down**, **Left**, **Right**) are used to select different controls on the screen. Alternatively the **Tab** key cycles through the controls on the screen which are enabled.

- Text fields are represented by a series of underscores (_). To enter data in a text field select it and begin entering data.
- Buttons are represented by labels which are enclosed within a pair of angle brackets (**<** and **>**). To activate a button ensure it is selected and press **Enter** or **Space**.
- Boolean options are represented by an asterisk (*****) or a space character enclosed within a pair of square brackets (**[** and **]**). When the value contained within the brackets is an asterisk then the option is set, otherwise it is not. To toggle a Boolean option on or off press **Space** while it is selected.

Report a bug

8.4.2. Installing the Hypervisor

Summary

There are two methods for installing Red Hat Enterprise Virtualization Hypervisors:

- Interactive installation.
- Unattended installation.

This section outlines the procedure for installing a Hypervisor interactively.

Procedure 8.10. Installing the Hypervisor Interactively

1. Use the prepared boot media to boot the machine on which the Hypervisor is to be installed.
2. Select **Install Hypervisor** and press **Enter** to begin the installation process.
3. The first screen that appears allows you to configure the appropriate keyboard layout for your locale. Use the arrow keys to highlight the appropriate option and press **Enter** to save your selection.

   Example 8.9. Keyboard Layout Configuration

   ```
   Keyboard Layout Selection

   Available Keyboard Layouts
   Swiss German (latin1)
   Turkish
   U.S. English
   U.S. International
   ...

   (Hit enter to select a layout)

   <Quit>     <Back>     <Continue>
   ```

4. The installation script automatically detects all disks attached to the system. This information is used to assist with selection of the boot and installation disks that the Hypervisor will use. Each entry displayed on these screens indicates the **Location**, **Device Name**, and **Size** of the disks.

a. **Boot Disk**

The first disk selection screen is used to select the disk from which the Hypervisor will boot. The Hypervisor's boot loader will be installed to the Master Boot Record (MBR) of the disk that is selected on this screen. The Hypervisor attempts to automatically detect the disks attached to the system and presents the list from which to choose the boot device. Alternatively, you can manually select a device by specifying a block device name using the **Other Device** option.

Important

The selected disk must be identified as a boot device and appear in the boot order either in the system's BIOS or in a pre-existing boot loader.

A. **Automatically Detected Device Selection**

i. Select the entry for the disk the Hypervisor is to boot from in the list and press **Enter**.

ii. Select the disk and press **Enter**. This action saves the boot device selection and starts the next step of installation.

B. **Manual Device Selection**

i. Select **Other device** and press **Enter**.

ii. When prompted to **Please select the disk to use for booting RHEV-H**, enter the name of the block device from which the Hypervisor should boot.

Example 8.10. Other Device Selection

```
Please select the disk to use for booting RHEV-H
/dev/sda
```

iii. Press **Enter**. This action saves the boot device selection and starts the next step of installation.

b. The disk or disks selected for installation will be those to which the Hypervisor itself is installed. The Hypervisor attempts to automatically detect the disks attached to the system and presents the list from which installation devices are chosen.

Warning

All data on the selected storage devices will be destroyed.

i. Select each disk on which the Hypervisor is to be installed and press **Space** to toggle it to enabled. Where other devices are to be used for installation, either solely or in addition to those which are listed automatically, use **Other Device**.

ii. Select the **Continue** button and press **Enter** to continue.

iii. Where the **Other Device** option was specified, a further prompt will appear. Enter the name of each additional block device to use for Hypervisor installation, separated by a comma. Once all required disks have been selected, select the **<Continue>** button and press **Enter**.

Example 8.11. Other Device Selection

```
Please enter one or more disks to use for installing
RHEV-H. Multiple devices can be separated by comma.
Device path:
/dev/mmcblk0,/dev/mmcblk1_______________
```

Once the installation disks have been selected, the next stage of the installation starts.

5. The next screen allows you to configure storage for the Hypervisor.

   a. Select or clear the **Fill disk with Data partition** check box. Clearing this text box displays a field showing the remaining space on the drive and allows you to specify the amount of space to be allocated to data storage.

   b. Enter the preferred values for **Swap**, **Config**, and **Logging**.

   c. If you selected the **Fill disk with Data partition** check box, the **Data** field is automatically set to **0**. If the check box was cleared, you can enter a whole number up to the value of the **Remaining Space** field. Entering a value of **-1** fills all remaining space.

6. The Hypervisor requires a password be set to protect local console access to the **admin** user. The installation script prompts you to enter the preferred password in both the **Password** and **Confirm Password** fields.

   Use a strong password. Strong passwords comprise a mix of uppercase, lowercase, numeric, and punctuation characters. They are six or more characters long and do not contain dictionary words.

   Once a strong password has been entered, select **<Install>** and press **Enter** to install the Hypervisor on the selected disks.

Result

Once installation is complete, the message **RHEV Hypervisor Installation Finished Successfully** will be displayed. Select the **<Reboot>** button and press **Enter** to reboot the system.

Note

Remove the boot media and change the boot device order to prevent the installation sequence restarting after the system reboots.

Note

Red Hat Enterprise Virtualization Hypervisors are able to use Storage Area Networks (SANs) and other network storage for storing virtualized guest images. Hypervisors can be installed on SANs, provided that the Host Bus Adapter (HBA) permits configuration as a boot device in BIOS.

Note

Hypervisors are able to use multipath devices for installation. Multipath is often used for SANs or other networked storage. Multipath is enabled by default at install time. Any block device which responds to **scsi_id** functions with multipath. Devices where this is not the case include USB storage and some older ATA disks.

Report a bug

8.5. Configuring the Hypervisor

8.5.1. Logging into the Hypervisor

Summary

The Hypervisor allows local console logins to facilitate post-installation configuration.

Procedure 8.11. Logging into the Hypervisor

1. Boot the Hypervisor. A login prompt appears:

```
Please login as 'admin' to configure the node
localhost login:
```

2. Enter the user name **admin** and press **Enter**.
3. Enter the password you set during Hypervisor installation and press **Enter**.

Result

You have successfully logged into the Hypervisor as the **admin** user.

Report a bug

8.5.2. Selecting Hypervisor Keyboard

Summary

The **Keyboard** screen allows you to configure the keyboard layout.

Procedure 8.12. Configuring the Hypervisor Keyboard Layout

1. Select a keyboard layout from the list provided.

```
Keyboard Layout Selection

Choose the Keyboard Layout you would like to apply to this system.

Current Active Keyboard Layout: U.S. English
Available Keyboard Layouts
Swiss German (latin1)
Turkish
U.S. English
U.S. International
Ukranian
...

<Save>
```

2. Select **Save** and press **Enter** to save the selection.

Result

You have successfully configured the keyboard layout.

Report a bug

8.5.3. Viewing Hypervisor Status

The status screen displays a brief overview of the current state of the Hypervisor. The information displayed consists of:

- The current status of the Hypervisor.
- The current status of networking.
- The destinations of logs and reports.
- The number of active virtual machines.

The status screen also provides a number of buttons for viewing further details on the Hypervisor and for changing the state of the Hypervisor:

- **<View Host Key>**: Displays the RSA host key fingerprint and host key of the Hypervisor.
- **<View CPU Details>**: Displays details on the CPU used by the Hypervisor such as the CPU name and type.
- **<Lock>**: Locks the Hypervisor. The user name and password must be entered to unlock the Hypervisor.
- **<Log Off>**: Logs off the current user.
- **<Restart>**: Restarts the Hypervisor.
- **<Power Off>**: Turns the Hypervisor off.

Report a bug

8.5.4. Configuring Hypervisor Network

8.5.4.1. Hypervisor Network Screen

The network screen is used to configure:

- The host name of the Hypervisor.
- The DNS servers to use.
- The NTP servers to use.
- The network interface to use.

The network screen also provides a number of buttons for testing and configuring network interfaces:

- **<Ping>**: Allows you to ping a given IP address by specifying the address to ping and number of times to ping that address.
- **<Create Bond>**: Allows you to create bonds between network interfaces.

See Also:

- Section 8.5.4.2, “Configuring Hypervisor Host Name”
- Section 8.5.4.3, “Configuring Hypervisor Domain Name Servers”
- Section 8.5.4.4, “Configuring Hypervisor Network Time Protocol”
- Section 8.5.4.5, “Configuring Hypervisor Network Interfaces”

Report a bug

8.5.4.2. Configuring Hypervisor Host Name

Summary

The Hypervisor host name is configurable.

Procedure 8.13. Configuring Hypervisor Host Name

1. Select the **Hostname** field on the **Network** screen and enter the new host name.
2. Select **<Save>** and press **Enter** to save changes to the host name.

Result

The host name has been updated.

Report a bug

8.5.4.3. Configuring Hypervisor Domain Name Servers

Summary

The Hypervisor supports the specification of one or more Domain Name Servers (DNS). These servers will be used when resolving host and domain names.

Procedure 8.14. Configuring Hypervisor Domain Name Servers

1. To set or change the primary DNS server, select the **DNS Server 1** field and enter the IP address of the new primary DNS server to use.
2. To set or change the secondary DNS server, select the **DNS Server 2** field and enter the IP address of the new secondary DNS server to use.
3. Select **<Save>** and press **Enter** to save changes to the DNS configuration.

Result

The primary and secondary DNS servers queried by the Hypervisor are updated.

Report a bug

8.5.4.4. Configuring Hypervisor Network Time Protocol

Summary

The Hypervisor supports the specification of one or more Network Time Protocol (NTP) servers with which the Hypervisor can synchronize its system clock. It is important that the Hypervisor is synchronized with the same time source as the Red Hat Enterprise Virtualization Manager. This ensures accurate time keeping across the Red Hat Enterprise Virtualization environment.

Procedure 8.15. Configuring Hypervisor Network Time Protocol

1. To set or change the primary NTP server, select the **NTP Server 1** field and enter the IP address or host name of the new primary NTP server to use.
2. To set or change the secondary NTP server, select the **NTP Server 2** field and enter the IP address or host name of the new secondary NTP server to use.
3. Select **<Save>** and press **Enter** to save changes to the NTP configuration.

Result

The primary and secondary NTP servers queried by the Hypervisor are updated.

Report a bug

8.5.4.5. Configuring Hypervisor Network Interfaces

Summary

For each network interface detected the Hypervisor will display the:

- **Device**
- **Status**
- **Model**
- **MAC Address**

At least one network interface must be configured before the Hypervisor can connect with the Red Hat Enterprise Virtualization Manager.

Procedure 8.16. Configuring Hypervisor Network Interfaces

1. **Device Identification**

   Select the network interface to be configured from the list and press **Enter**.

   When it is unclear which physical device an entry in the list refers to, the Hypervisor can blink the network traffic lights on the physical device to assist with identification. To use this facility, select the entry from the list, select the **<Flash Lights to Identify>** button and press **Enter**. Take note of which physical device's lights start blinking. The configuration screen for the selected device will be displayed.

2. **IPv4 Settings**

   The Hypervisor supports both dynamic (DHCP) and static IPv4 network configuration.

   A. **Dynamic (DHCP) Network Configuration**

   Dynamic network configuration allows the Hypervisor to be dynamically assigned an IP address via **DHCP**. To enable dynamic IPv4 network configuration, select the **DHCP** option under **IPv4 Settings** and press **Space** to toggle it to enabled.

   B. **Static Network Configuration**

   Static network configuration allows the Hypervisor to be manually assigned an IP address. To enable static IPv4 network configuration select the **Static** option under **IPv4 Settings** and press **Space** to toggle it to enabled.

   Selection of the **Static** option enables the **IP Address**, **Netmask**, and **Gateway** fields. The **IP Address**, **Netmask**, and **Gateway** fields must be populated to complete static network configuration.

   In particular it is necessary that:

   - the **IP Address** is not already in use on the network,
   - the **Netmask** matches that used by other machines on the network, and
   - the **Gateway** matches that used by other machines on the network.

   Where it is not clear what value should be used for the **IP Address**, **Netmask**, or **Gateway** field consult the network's administrator or consider a dynamic configuration.

   Example 8.12. Static IPv4 Networking Configuration

   ```
   IPv4 Settings
   ( ) Disabled     ( ) DHCP     (*) Static
   IP Address: 192.168.122.100_  Netmask: 255.255.255.0___
   Gateway     192.168.1.1_____
   ```

3. **IPv6 Settings**

   The Red Hat Enterprise Virtualization Manager does not currently support IPv6 networking. IPv6 networking must remain set to **Disabled**.

4. **VLAN Configuration**

If VLAN support is required, populate the **VLAN ID** field with the VLAN identifier for the selected device.

5. **Save Network Configuration**

   Once all networking options for the selected device have been set, the configuration must be saved.

   a. Select the **<Save>** button and press **Enter** to save the network configuration.

   b. A screen showing the progress of configuration displays. Once configuration is complete, press the **Enter** key to close the window.

Result

The network screen is displayed. The device is listed as **Configured**.

Report a bug

8.5.5. Configuring Hypervisor Security

Summary

The **Security** screen is used to change the **admin** password for both local and remote access. SSH password authentication is also enabled or disabled via this screen.

Procedure 8.17. Configuring Hypervisor Security

1. **Enable SSH Password Authentication**

   To enable SSH password authentication for remote access, select the **Enable ssh password authentication** option and press **Space** to toggle it to enabled.

2. **Change admin Password**

   a. Enter the desired **admin** password in the **Password** field. You should use a strong password.

   Strong passwords contain a mix of uppercase, lowercase, numeric and punctuation characters. They are six or more characters long and do not contain dictionary words.

   b. Enter the desired **admin** password in the **Confirm Password** field. Ensure the value entered in the **Confirm Password** field matches the value entered in the **Password** field exactly. Where this is not the case, an error message will be displayed to indicate that the two values are different.

3. Select **<Save>** and press **Enter** to save the security configuration.

Result

The security configuration has been updated.

Report a bug

8.5.6. Configuring Hypervisor Simple Network Management Protocol

Summary

The SNMP screen allows you to enable and configure a password for Simple Network Management Protocol.

```
Enable SNMP       [ ]

SNMP Password
Password:          _________________
Confirm Password:  _________________


<Save>     <Reset>
```

Procedure 8.18. Configuring Hypervisor Simple Network Management Protocol

1. Select the **Enable SNMP** field.
2. Press **Space** to toggle between enabling SNMP and disabling SNMP. By default, SNMP is disabled.
3. Enter the preferred SNMP **Password** for the Hypervisor.
4. Enter the preferred SNMP password again in the **Confirm Password** field.
5. Select **<Save>** and press **Enter** to save your changes.

Result

SNMP is now configured.

Report a bug

8.5.7. Configuring Hypervisor Common Information Model

Summary

Enabling and configuring Common Information Model (CIM) allows you to attach the Hypervisor to your pre-existing CIM management infrastructure and monitor virtual machines that are run on the Hypervisor.

Procedure 8.19. Configuring Hypervisor Common Information Model

1. Select the **Enable CIM** field.

   ```
   Enable CIM     [ ]
   ```

2. Enter a password in the **Password** field. This is the password that you will use to access the Hypervisor using CIM.
3. Enter the password again in the **Confirm Password** field.
4. Select the **Save** button and press **Enter** to save your changes.

Result

You have configured the Hypervisor to accept CIM connections authenticated using a password. Use this password when adding the Hypervisor to your Common Information Model Object Manager (CIMOM).

Report a bug

8.5.8. Configuring Logging

Summary

The Hypervisor creates and updates a number of log files. The **Logging** screen allows configuration of a daemon to automatically export these log files to a remote server.

Procedure 8.20. Configuring Hypervisor Logging

1. **logrotate Configuration**

   The **logrotate** utility simplifies the administration of log files. The Hypervisor uses **logrotate** to rotate logs when they reach a certain file size.

   Log rotation involves renaming the current logs and starting new ones in their place. The **Logrotate Max Log Size** value set on the **Logging** screen is used to determine when a log will be rotated.

   Enter the **Logrotate Max Log Size** in kilobytes. The default maximum log size is 1024 kilobytes.

2. **rsyslog Configuration**

   The **rsyslog** utility is a multithreaded syslog daemon. The Hypervisor is able to use **rsyslog** to transmit log files over the network to a remote syslog daemon. For information on setting up the remote syslog daemon, see the *Red Hat Enterprise Linux Deployment Guide*.

   a. Enter the remote **rsyslog** server address in the **Server Address** field.

   b. Enter the remote **rsyslog** server port in the **Server Port** field. The default port is **514**.

3. **netconsole Configuration**

   The **netconsole** module allows kernel messages to be sent to a remote machine. The Hypervisor uses **netconsole** to transmit kernel messages over the network.

   a. Enter the **Server Address**.

   b. Enter the **Server Port**. The default port is **6666**.

4. **Save Configuration**

   To save the logging configuration, select **<Save>** and press **Enter**.

Result

The logging configuration has been updated and logs will be exported to the remote **rsyslog** server specified.

Report a bug

8.5.9. Configuring the Hypervisor for Red Hat Network

Summary

Guests running on the Hypervisor may need to consume Red Hat Enterprise Linux virtualization entitlements. In this case, the Hypervisor must be registered to Red Hat Network, a Satellite server, or Subscription Asset Manager. The Hypervisor can connect to these services via a HTTP proxy where one is in use.

Note that when new versions of the Hypervisor itself become available, they are installed from the Red Hat Enterprise Virtualization Manager and not Red Hat Network.

Procedure 8.21. Configuring Hypervisor for Red Hat Network

1. **Authentication**

   Enter your Red Hat Network user name in the **Login** field.

   Enter your Red Hat Network password in the **Password** field.

2. **Profile Name**

   Enter the profile name to be used for the system in the **Profile Name** field. This is the name that the system will appear under when viewed in Red Hat Network.

3. **Update Source**

   The Hypervisor can register directly to Red Hat Network or, if available, a Satellite installation or a Subscription Asset Manager.

   A. **To Connect Directly to RHN**

   Select the **RHN** option and press **Space** to toggle it to enabled. The **RHN URL** and **CA URL** values do not need to be provided.

   Example 8.13. Red Hat Network Configuration

   ```
   (X) RHN     ( ) Satellite     ( ) SAM
   RHN URL:
   ____________________________________________________________________
   _
   CA URL:
   ____________________________________________________________________
   _
   ```

   B. **To Connect via Satellite**

   a. Select the **Satellite** option and press **Space** to toggle it to enabled.

   b. Enter the URL of the Satellite server in the **RHN URL** field.

   c. Enter the URL of the certificate authority for the Satellite server in the **CA URL** field.

   Example 8.14. Satellite Configuration

```
( ) RHN     (X) Satellite     ( ) SAM
RHN URL:   https://your-
satellite.example.com_____________________________
CA URL:    https://your-satellite.example.com/pub/RHN-ORG-
TRUSTED-SSL-CERT
```

C. **To Connect via Subscription Asset Manager**

a. Select the **Subscription Asset Manager** option and press **Space** to toggle it to enabled.

b. Enter the URL of the Subscription Asset Manager server in the **RHN URL** field.

c. Enter the URL of the certificate authority for the Subscription Asset Manager server in the **CA URL** field.

Example 8.15. Subscription Asset Manager Configuration

```
( ) RHN     ( ) Satellite     (X) SAM
URL:  https://subscription-asset-
manager.example.com_____________________________
CA :  https://subscription-asset-manager.example.com/pub/RHN-
ORG-TRUSTED-SSL-CERT
```

4. **HTTP Proxy**

Where a HTTP proxy is in use the details to connect to it must be provided. To connect to Red Hat Network or a Satellite server via a proxy you must enter:

- The network address of the proxy **Server**.
- The **Port** to connect to the proxy on.
- Optionally, the **Username** and **Password** to use to connect to the proxy.

In environments where a HTTP proxy is not in use, you can ignore this step.

Example 8.16. HTTP Proxy Configuration

```
HTTP Proxy Configuration

Server:   proxy.example.com__
Port:     8080_____________
Username: puser____________
Password: ******____________
```

5. **Save Configuration**

To save the configuration the user must select **<Save>** and press **Enter**.

Result

The **Red Hat Network** configuration of the Hypervisor has been updated.

Report a bug

8.5.10. Configuring Hypervisor Kernel Dumps

Summary

The Hypervisor supports the export of kernel dumps in the event of a system failure using the **kdump** facility. The **kdump** files can be delivered using NFS or SSH so that they can be analyzed at a later date. The Kdump screen allows you to configure this facility.

Procedure 8.22. Configuring Hypervisor Kernel Dumps

1. Crash dumps generated by **kdump** are exported over NFS or SSH. Select the preferred transfer method and press **Space** to enable it.

   For the selected export method, a location to which the **kdump** files are to be exported must also be specified.

   a. **NFS Location**

   Set the NFS location to which crash logs are to be exported in the **NFS Location** field. The **NFS Location** must be the full NFS path which includes fully qualified domain name and directory path.

   > **Example 8.17. NFS Location**
   >
   > **example.redhat.com:/var/crash**

   b. **SSH Location**

   Set the SSH location to which crash logs are to be exported in the **SSH Location** field. The **SSH Location** must be the full SSH login which includes the fully qualified domain name and user name.

   > **Example 8.18. SSH Location**
   >
   > **root@example.redhat.com**

2. **Save Configuration**

   To save the configuration, select **<Save>** and press **Enter**.

Result

The kernel dump configuration has been updated and kernel dumps will be exported to the specified remote servers.

Report a bug

8.5.11. Configuring Hypervisor Remote Storage

Summary

The Hypervisor supports the use of a remote iSCSI initiator for storage. The iSCSI initiator to use is set from the **Remote Storage** screen.

Procedure 8.23. Configuring Hypervisor Remote Storage

1. **iSCSI Initiator Name**

   Enter the initiator name in the **iSCSI Initiator Name** field.

   > **Example 8.19. iSCSI Initiator Name**
   >
   > **iqn.1994-05.com.redhat:5189835eeb40**

2. **Save Configuration**

   To save the configuration the user must select **<Save>** and press **Enter**.

Result

The **Remote Storage** configuration has been updated.

Report a bug

8.6. Attaching the Hypervisor to the Red Hat Enterprise Virtualization Manager

8.6.1. Configuring Hypervisor Management Server

Summary

You can attach the Hypervisor to the Red Hat Enterprise Virtualization Manager immediately if the address of the Manager is available. If the Manager has not yet been installed, you must instead set a password. This allows the Hypervisor to be added from the Administration Portal once the Manager has been installed. Both modes of configuration are supported from the **oVirt Engine** screen in the Hypervisor user interface.

Important

Setting a password on the **oVirt Engine** configuration screen sets the **root** password on the Hypervisor and enables SSH password authentication. Once the Hypervisor has successfully been added to the Manager, disabling SSH password authentication is recommended.

Procedure 8.24. Configuring a Hypervisor Management Server

1. A. **Configuration Using a Management Server Address**

   a. Enter the IP address or fully qualified domain name of the Manager in the **Management Server** field.

b. Enter the management server port in the **Management Server Port** field. The default value is **443**. If a different port was selected during Red Hat Enterprise Virtualization Manager installation, specify it here, replacing the default value.

c. Select the **Retrieve Certificate** option to verify that the fingerprint of the certificate retrieved from the specified management server is correct. The value that the certificate fingerprint is compared against is returned at the end of Red Hat Enterprise Virtualization Manager installation.

d. Leave the **Password** and **Confirm Password** fields blank. These fields are not required if the address of the management server is known.

B. **Configuration Using a Password**

a. Enter a password in the **Password** field. It is recommended that you use a strong password. Strong passwords contain a mix of uppercase, lowercase, numeric and punctuation characters. They are six or more characters long and do not contain dictionary words.

b. Re-enter the password in the **Confirm Password** field.

c. Leave the **Management Server** and **Management Server Port** fields blank. As long as a password is set, allowing the Hypervisor to be added to the Manager later, these fields are not required.

2. **Save Configuration**

To save the configuration select **<Save>** and press **Enter**.

Result

The **oVirt Engine** configuration has been updated.

See Also:

- Section 8.6.3, “Approving a Hypervisor”

Report a bug

8.6.2. Using the Hypervisor

If the Hypervisor was configured with the address of the Red Hat Enterprise Virtualization Manager, the Hypervisor reboots and is automatically registered with the Manager. The Red Hat Enterprise Virtualization Manager interface displays the Hypervisor under the **Hosts** tab. To prepare the Hypervisor for use, it must be approved using Red Hat Enterprise Virtualization Manager.

If the Hypervisor was configured without the address of the Red Hat Enterprise Virtualization Manager, it must be added manually. To add the Hypervisor manually, you must have both the IP address of the machine upon which it was installed and the password that was set on the **oVirt Engine** screen during configuration.

Report a bug

8.6.3. Approving a Hypervisor

Summary

It is not possible to run virtual machines on a Hypervisor until the addition of it to the environment has been approved in Red Hat Enterprise Virtualization Manager.

Procedure 8.25. Approving a Hypervisor

1. Log in to the Red Hat Enterprise Virtualization Manager Administration Portal.
2. From the **Hosts** tab, click on the host to be approved. The host should currently be listed with the status of **Pending Approval**.
3. Click the **Approve** button. The **Edit and Approve Hosts** dialog displays. You can use the dialog to set a name for the host, fetch its SSH fingerprint before approving it, and configure power management, where the host has a supported power management card. For information on power management configuration, see the *Power Management* chapter of the *Red Hat Enterprise Virtualization Administration Guide*.
4. Click **OK**. If you have not configured power management you will be prompted to confirm that you wish to proceed without doing so, click **OK**.

Result

The status in the **Hosts** tab changes to **Installing**, after a brief delay the host status changes to **Up**.

See Also:

- Section 9.5, “Explanation of Settings and Controls in the New Host and Edit Host Windows”

Report a bug

Chapter 9. Installing Red Hat Enterprise Linux Hosts

9.1. Red Hat Enterprise Linux Hosts

You can use a standard Red Hat Enterprise Linux 6 installation on capable hardware as a host. Red Hat Enterprise Virtualization supports hosts running Red Hat Enterprise Linux 6 Server AMD64/Intel 64 version.

Adding a host can take some time, as the following steps are completed by the platform: virtualization checks, installation of packages, creation of bridge and a reboot of the host. Use the Details pane to monitor the hand-shake process as the host and management system establish a connection.

See Also:

- Section 9.2, “Host Compatibility Matrix”

Report a bug

9.2. Host Compatibility Matrix

Red Hat Enterprise Linux Version	Red Hat Enterprise Virtualization 3.3 clusters in 3.0 compatibility mode	Red Hat Enterprise Virtualization 3.3 clusters in 3.1 compatibility mode	Red Hat Enterprise Virtualization 3.3 clusters in 3.2 compatibility mode	Red Hat Enterprise Virtualization 3.3 clusters
6.2	Supported	Unsupported	Unsupported	Unsupported
6.3	Supported	Supported	Unsupported	Unsupported
6.4	Supported	Supported	Supported	Unsupported
6.5	Supported	Supported	Supported	Supported

Report a bug

9.3. Preparing a Red Hat Enterprise Linux Host

9.3.1. Installing Red Hat Enterprise Linux

Summary

You must install Red Hat Enterprise Linux 6.5 Server on a system to use it as a virtualization host in a Red Hat Enterprise Virtualization 3.3 environment.

Procedure 9.1. Installing Red Hat Enterprise Linux

1. **Download and Install Red Hat Enterprise Linux 6.5 Server**

   Download and Install Red Hat Enterprise Linux 6.5 Server on the target virtualization host, referring to the *Red Hat Enterprise Linux 6 Installation Guide* for detailed instructions. Only the *Base* package group is required to use the virtualization host in a Red Hat Enterprise Virtualization environment.

Important

If you intend to use directory services for authentication on the Red Hat Enterprise Linux host then you must ensure that the authentication files required by the **useradd** command are locally accessible. The *vdsm* package, which provides software that is required for successful connection to Red Hat Enterprise Virtualization Manager, will not install correctly if these files are not locally accessible.

2. **Ensure Network Connectivity**

   Following successful installation of Red Hat Enterprise Linux 6.5 Server, ensure that there is network connectivity between your new Red Hat Enterprise Linux host and the system on which your Red Hat Enterprise Virtualization Manager is installed.

   a. Attempt to ping the Manager:

   ```
   # ping address of manager
   ```

   b. A. If the Manager can successfully be contacted, this displays:

   ```
   ping manager.example.redhat.com
   PING manager.example.redhat.com (192.168.0.1) 56(84) bytes
   of data.
   64 bytes from 192.168.0.1: icmp_seq=1 ttl=64 time=0.415 ms
   64 bytes from 192.168.0.1: icmp_seq=2 ttl=64 time=0.419 ms
   64 bytes from 192.168.0.1: icmp_seq=3 ttl=64 time=1.41 ms
   64 bytes from 192.168.0.1: icmp_seq=4 ttl=64 time=0.487 ms
   64 bytes from 192.168.0.1: icmp_seq=5 ttl=64 time=0.409 ms
   64 bytes from 192.168.0.1: icmp_seq=6 ttl=64 time=0.372 ms
   64 bytes from 192.168.0.1: icmp_seq=7 ttl=64 time=0.464 ms

   --- manager.example.redhat.com ping statistics ---
   7 packets transmitted, 7 received, 0% packet loss, time
   6267ms
   ```

   B. If the Manager cannot be contacted, this displays:

   ```
   ping: unknown host manager.usersys.redhat.com
   ```

   You must configure the network so that the host can contact the Manager. First, disable **NetworkManager**. Then configure the networking scripts so that the host will acquire an ip address on boot.

   i. Disable **NetworkManager**.

   ```
   # service NetworkManager stop
   # chkconfig NetworkManager disable
   ```

   ii. Edit **/etc/sysconfig/network-scripts/ifcfg-eth0**. Find this line:

   ```
   ONBOOT=no
   ```

Change that line to this:

```
ONBOOT=yes
```

iii. Reboot the host machine.

iv. Ping the Manager again:

```
# ping address of manager
```

If the host still cannot contact the Manager, it is possible that your host machine is not acquiring an IP address from DHCP. Confirm that DHCP is properly configured and that your host machine is properly configured to acquire an IP address from DHCP.

If the Manager can successfully be contacted, this displays:

```
ping manager.example.redhat.com
PING manager.example.redhat.com (192.168.0.1) 56(84)
bytes of data.
64 bytes from 192.168.0.1: icmp_seq=1 ttl=64
time=0.415 ms
64 bytes from 192.168.0.1: icmp_seq=2 ttl=64
time=0.419 ms
64 bytes from 192.168.0.1: icmp_seq=3 ttl=64 time=1.41
ms
64 bytes from 192.168.0.1: icmp_seq=4 ttl=64
time=0.487 ms
64 bytes from 192.168.0.1: icmp_seq=5 ttl=64
time=0.409 ms
64 bytes from 192.168.0.1: icmp_seq=6 ttl=64
time=0.372 ms
64 bytes from 192.168.0.1: icmp_seq=7 ttl=64
time=0.464 ms

--- manager.example.redhat.com ping statistics ---
7 packets transmitted, 7 received, 0% packet loss,
time 6267ms
```

Result

You have installed Red Hat Enterprise Linux 6.5 Server. You must however complete additional configuration tasks before adding the virtualization host to your Red Hat Enterprise Virtualization environment.

Next Step in Preparing a Red Hat Enterprise Linux Host

- Section 9.3.2, “Subscribing to Required Channels Using Subscription Manager”

Report a bug

9.3.2. Subscribing to Required Channels Using Subscription Manager

Previous Step in Preparing a Red Hat Enterprise Linux Host

- Section 9.3.1, “Installing Red Hat Enterprise Linux”

Summary

To be used as a virtualization host, a Red Hat Enterprise Linux host must be registered and subscribed to a number of entitlements using either Subscription Manager or RHN Classic. You must follow the steps in this procedure to register and subscribe using Subscription Manager. Completion of this procedure will mean that you have:

- Registered the virtualization host to Red Hat Network using Subscription Manager.
- Attached the **Red Hat Enterprise Linux Server** entitlement to the virtualization host.
- Attached the **Red Hat Enterprise Virtualization** entitlement to the virtualization host.

Do *not* follow the steps in this procedure to register and subscribe using RHN Classic.

Procedure 9.2. Subscribing to Required Channels using Subscription Manager

1. **Register**

   Run the **subscription-manager** command with the ***register*** parameter to register the system with Red Hat Network. To complete registration successfully, you will need to supply your Red Hat Network **Username** and **Password** when prompted.

   ```
   # subscription-manager register
   ```

2. **Identify Available Entitlement Pools**

   To attach the correct entitlements to the system, you must first locate the identifiers for the required entitlement pools. Use the ***list*** action of the **subscription-manager** to find these.

   To identify available subscription pools for **Red Hat Enterprise Linux Server**, use the command:

   ```
   # subscription-manager list --available | grep -A8 "Red Hat
   Enterprise Linux Server"
   ```

   To identify available subscription pools for **Red Hat Enterprise Virtualization**, use the command:

   ```
   # subscription-manager list --available | grep -A8 "Red Hat
   Enterprise Virtualization"
   ```

3. **Attach Entitlements to the System**

   Using the pool identifiers you located in the previous step, attach the **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** entitlements to the system. Use the ***attach*** parameter of the **subscription-manager** command, replacing ***[POOLID]*** with each of the pool identifiers:

   ```
   # subscription-manager attach --pool=[POOLID]
   ```

4. **Enable the Red Hat Enterprise Virtualization Management Agents Repository**

Run the following command to enable the Red Hat Enterprise Virtualization Management Agents (RPMs) repository:

```
# subscription-manager repos --enable=rhel-6-server-rhev-mgmt-agent-rpms
```

Result

You have registered the virtualization host to Red Hat Network and attached the required entitlements using Subscription Manager.

Next Step in Preparing a Red Hat Enterprise Linux Host

- Section 9.3.3, “Subscribing to Required Channels Using RHN Classic”

Report a bug

9.3.3. Subscribing to Required Channels Using RHN Classic

Previous Step in Preparing a Red Hat Enterprise Linux Host

- Section 9.3.2, “Subscribing to Required Channels Using Subscription Manager”

Summary

To be used as a virtualization host, a Red Hat Enterprise Linux host must be registered and subscribed to a number of entitlements using either Subscription Manager or RHN Classic. You must follow the steps in this procedure if you wish to register and subscribe using RHN Classic. Completion of this procedure will mean that you have:

- Registered the virtualization host to Red Hat Network using RHN Classic.
- Subscribed the virtualization host to the **Red Hat Enterprise Linux Server (v. 6 for 64-bit AMD64 / Intel64)** channel.
- Subscribed the virtualization host to the **Red Hat Enterprise Virt Management Agent (v 6 x86_64)** channel.

Do *not* follow the steps in this procedure if you wish to register and subscribe using Subscription Manager.

Procedure 9.3. Subscribing to Required Channels using RHN Classic

1. **Register**

   If the machine has not already been registered with Red Hat Network, run the **rhn_register** command as **root** to register it. To complete registration successfully you will need to supply your Red Hat Network **Username** and **Password**. Follow the prompts displayed by **rhn_register** to complete registration of the system.

   ```
   # rhn_register
   ```

2. **Subscribe to channels**

   You must subscribe the system to the required channels using either the web interface to Red Hat Network or the command line **rhn-channel** command.

A. **Using the Web Interface to Red Hat Network**

To add a channel subscription to a system from the web interface:

a. Log on to Red Hat Network (http://rhn.redhat.com).

b. Move the mouse cursor over the **Subscriptions** link at the top of the screen, and then click the **Registered Systems** link in the menu that appears.

c. Select the system to which you are adding channels from the list presented on the screen, by clicking the name of the system.

d. Click **Alter Channel Subscriptions** in the **Subscribed Channels** section of the screen.

e. Select the channels to be added from the list presented on the screen. To use the virtualization host in a Red Hat Enterprise Virtualization environment you must select:

- **Red Hat Enterprise Linux Server (v. 6 for 64-bit x86_64)**; and
- **Red Hat Enterprise Virt Management Agent (v 6 x86_64)**.

f. Click the **Change Subscription** button to finalize the change.

B. **Using the rhn-channel command**

Run the **rhn-channel** command to subscribe the virtualization host to each of the required channels. The commands that need to be run are:

```
# rhn-channel --add --channel=rhel-x86_64-server-6
# rhn-channel --add --channel=rhel-x86_64-rhev-mgmt-agent-6
```

Important

If you are not the administrator for the machine as defined in Red Hat Network, or the machine is not registered to Red Hat Network, then use of the **rhn-channel** command will result in an error:

```
Error communicating with server. The message was:Error Class
Code: 37
Error Class Info: You are not allowed to perform
administrative tasks on this system.
Explanation:
     An error has occurred while processing your request. If
this problem
     persists please enter a bug report at
bugzilla.redhat.com.
     If you choose to submit the bug report, please be sure
to include
     details of what you were trying to do when this error
occurred and
     details on how to reproduce this problem.
```

If you encounter this error when using **rhn-channel** then to add the channel to the system you must use the web user interface instead.

Result

You have registered the virtualization host to Red Hat Network and subscribed to the required entitlements using RHN Classic.

Next Step in Preparing a Red Hat Enterprise Linux Host

- Section 9.3.4, “Configuring Virtualization Host Firewall”

Report a bug

9.3.4. Configuring Virtualization Host Firewall

Previous Step in Preparing a Red Hat Enterprise Linux Host

- Section 9.3.3, “Subscribing to Required Channels Using RHN Classic”

Summary

Red Hat Enterprise Virtualization requires that a number of network ports be open to support virtual machines and remote management of the virtualization host from the Red Hat Enterprise Virtualization Manager. You must follow this procedure to open the required network ports before attempting to add the virtualization host to the Manager.

Procedure 9.4. Configuring Virtualization Host Firewall

The following steps configure the default firewall in Red Hat Enterprise Linux, **iptables**, to allow traffic on the required network ports. These steps replace any existing firewall configuration on your host with one containing only the required by Red Hat Enterprise Virtualization. If you have existing

firewall rules with which this configuration must be merged then you must do so by manually editing the rules defined in the **iptables** configuration file, **/etc/sysconfig/iptables**.

All commands in this procedure must be run as the **root** user.

1. **Remove existing firewall rules from configuration**

   Remove any existing firewall rules using the ***--flush*** parameter to the **iptables** command.

   ```
   # iptables --flush
   ```

2. **Add new firewall rules to configuration**

   Add the new firewall rules, required by Red Hat Enterprise Virtualization, using the ***--append*** parameter to the **iptables** command. The prompt character (#) has been intentionally omitted from this list of commands to allow easy copying of the content to a script file or command prompt.

   ```
   iptables --append INPUT -m state --state ESTABLISHED,RELATED -j
   ACCEPT
   iptables --append INPUT -p icmp -j ACCEPT
   iptables --append INPUT -i lo -j ACCEPT
   iptables --append INPUT -p tcp --dport 22 -j ACCEPT
   iptables --append INPUT -p tcp --dport 16514 -j ACCEPT
   iptables --append INPUT -p tcp --dport 54321 -j ACCEPT
   iptables --append INPUT -p tcp -m multiport --dports 5634:6166 -j
   ACCEPT
   iptables --append INPUT -p tcp -m multiport --dports 49152:49216 -j
   ACCEPT
   iptables --append INPUT -j REJECT --reject-with icmp-host-
   prohibited
   iptables --append FORWARD -m physdev ! --physdev-is-bridged -j
   REJECT \
   --reject-with icmp-host-prohibited
   ```

   Note

   The provided **iptables** commands add firewall rules to accept network traffic on a number of ports. These include:
   - port **22** for **SSH**,
   - ports **5634** to **6166** for guest console connections,
   - port **16514** for **libvirt** virtual machine migration traffic,
   - ports **49152** to **49216** for VDSM virtual machine migration traffic, and
   - port **54321** for the Red Hat Enterprise Virtualization Manager.

3. **Save the updated firewall configuration**

   Save the updated firewall configuration script using the ***save*** to the **iptables** initialization script.

   ```
   # service iptables save
   ```

4. **Enable iptables service**

   Ensure that the **iptables** service is configured to start on boot and has been restarted, or started for the first time if it was not already running.

```
# chkconfig iptables on
# service iptables restart
```

Result

You have configured the virtualization host's firewall to allow the network traffic required by Red Hat Enterprise Virtualization.

Next Step in Preparing a Red Hat Enterprise Linux Host

- Section 9.3.5, “Configuring Virtualization Host sudo”

See Also:

- Section 2.3.5.1, “Red Hat Enterprise Virtualization Manager Firewall Requirements”

Report a bug

9.3.5. Configuring Virtualization Host sudo

Previous Step in Preparing a Red Hat Enterprise Linux Host

- Section 9.3.4, “Configuring Virtualization Host Firewall”

Summary

The Red Hat Enterprise Virtualization Manager makes use of **sudo** to perform operations as the **root** on the host. The default Red Hat Enterprise Linux configuration, stored in **/etc/sudoers**, contains values that allow this. If this file has been modified since Red Hat Enterprise Linux installation these values may have been removed. This procedure provides steps to verify that the required entry still exists in the configuration, and add the required entry if it is not present.

Procedure 9.5. Configuring Virtualization Host sudo

1. **Log in**

   Log in to the virtualization host as the **root** user.

2. **Run visudo**

   Run the **visudo** command to open the **/etc/sudoers**

```
# visudo
```

3. **Edit sudoers file**

   Read the configuration file, and verify that it contains these lines:

```
# Allow root to run any commands anywhere
root    ALL=(ALL)   ALL
```

If the file does not contain these lines, add them and save the file using the VIM **:w** command.

4. **Exit editor**

   Exit **visudo** using the VIM **:q** command.

Result

You have configured **sudo** to allow use by the **root** user.

Next Step in Preparing a Red Hat Enterprise Linux Host

- Section 9.3.6, “Configuring Virtualization Host SSH”

Report a bug

9.3.6. Configuring Virtualization Host SSH

Previous Step in Preparing a Red Hat Enterprise Linux Host

- Section 9.3.5, “Configuring Virtualization Host sudo”

Summary

The Red Hat Enterprise Virtualization Manager accesses virtualization hosts via SSH. To do this it logs in as the **root** user using an encrypted key for authentication. You must follow this procedure to ensure that SSH is configured to allow this.

Warning

The first time the Red Hat Enterprise Virtualization Manager is connected to the host it will install an authentication key. In the process it will overwrite any existing keys contained in the **/root/.ssh/authorized_keys** file.

Procedure 9.6. Configuring virtualization host SSH

All commands in this procedure must be run as the **root** user.

1. **Install the SSH server (*openssh-server*)**

   Install the *openssh-server* package using **yum**.

   ```
   # yum install openssh-server
   ```

2. **Edit SSH server configuration**

   Open the SSH server configuration, **/etc/ssh/sshd_config**, in a text editor. Search for the ***PermitRootLogin***.

   A. If ***PermitRootLogin*** is set to **yes**, or is not set at all, no further action is required.

   B. If ***PermitRootLogin*** is set to **no**, then you must change it to **yes**.

   Save any changes that you have made to the file, and exit the text editor.

3. **Enable the SSH server**

   Configure the SSH server to start at system boot using the **chkconfig** command.

   ```
   # chkconfig --level 345 sshd on
   ```

4. **Start the SSH server**

   Start the SSH, or restart it if it is already running, using the **service** command.

   ```
   # service sshd restart
   ```

Result

You have configured the virtualization host to allow **root** access over SSH.

Report a bug

9.4. Adding a Red Hat Enterprise Linux Host

Summary

A Red Hat Enterprise Linux host is based on a standard "basic" installation of Red Hat Enterprise Linux. The physical host must be set up before you can add it the Red Hat Enterprise Virtualization environment.

The Red Hat Enterprise Virtualization Manager logs into the host to perform virtualization capability checks, install packages, create a network bridge, and reboot the host. The process of adding a new host can take up to 10 minutes.

Procedure 9.7. Adding a Red Hat Enterprise Linux Host

1. Click the **Hosts** resource tab to list the hosts in the results list.
2. Click **New** to open the **New Host** window.
3. Use the drop-down menus to select the **Data Center** and **Host Cluster** for the new host.
4. Enter the **Name**, **Address**, and **SSH Port** of the new host.
5. Select an authentication method to use with the host.
   A. Enter the root user's password to use password authentication.
   B. Copy the key displayed in the **SSH PublicKey** field to **/root/.ssh/authorized_keys** on the host to use public key authentication.
6. You have now completed the mandatory steps to add a Red Hat Enterprise Linux host. Click the **Advanced Parameters** button to expand the advanced host settings.
   a. Optionally disable automatic firewall configuration.
   b. Optionally add a host SSH fingerprint to increase security. You can add it manually, or fetch it automatically.

7. You can configure the **Power Management** and **SPM** using the applicable tabs now; however, as these are not fundamental to adding a Red Hat Enterprise Linux host, they are not covered in this procedure.

8. Click **OK** to add the host and close the window.

Result

The new host displays in the list of hosts with a status of **Installing**. Once installation is complete, the status will update to **Reboot**. The host must be activated for the status to change to **Up**.

Note

You can view the progress of the installation in the details pane.

See Also:

- Section 9.5, “Explanation of Settings and Controls in the New Host and Edit Host Windows”

Report a bug

9.5. Explanation of Settings and Controls in the New Host and Edit Host Windows

9.5.1. Host General Settings Explained

These settings apply when editing the details of a host or adding new Red Hat Enterprise Linux hosts and Foreman host provider hosts.

The **General** settings table contains the information required on the **General** tab of the **New Host** or **Edit Host** window.

Table 9.1. General settings

Field Name	Description
Data Center	The data center to which the host belongs. Red Hat Enterprise Virtualization Hypervisor hosts can not be added to Gluster-enabled clusters.
Host Cluster	The cluster to which the host belongs.

Field Name	Description
Use External Providers	Select or clear this check box to view or hide options for adding hosts provided by external providers. Upon selection, a drop-down list of external providers that have been added to the Manager displays. The following options are also available: - **Provider search filter** - A text field that allows you to search for hosts provided by the selected external provider. This option is provider-specific; see provider documentation for details on forming search queries for specific providers. Leave this field blank to view all available hosts. - **External Hosts** - A drop-down list that is populated with the name of hosts provided by the selected external provider. The entries in this list are filtered in accordance with any search queries that have been input in the **Provider search query** field.
Name	The name of the cluster. This text field has a 40-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
Comment	A field for adding plain text, human-readable comments regarding the host.
Address	The IP address, or resolvable hostname of the host.
Root password	The password of the host's root user. This can only be given when you add the host, it cannot be edited afterwards.
SSH PublicKey	Copy the contents in the text box to the **/root/.known_hosts** file on the host if you'd like to use the Manager's ssh key instead of using a password to authenticate with the host.
Automatically configure host firewall	When adding a new host, the Manager can open the required ports on the host's firewall. This is enabled by default. This is an **Advanced Parameter**.
SSH Fingerprint	You can **fetch** the host's SSH fingerprint, and compare it with the fingerprint you expect the host to return, ensuring that they match. This is an **Advanced Parameter**.

Report a bug

9.5.2. Host Power Management Settings Explained

The **Power Management** settings table contains the information required on the **Power Management** tab of the **New Host** or **Edit Host** windows.

Table 9.2. Power Management Settings

Field Name	Description
Primary/ Secondary	Prior to Red Hat Enterprise Virtualization 3.2, a host with power management configured only recognized one fencing agent. Fencing agents configured on version 3.1 and earlier, and single agents, are treated as primary agents. The secondary option is valid when a second agent is defined.
Concurrent	Valid when there are two fencing agents, for example for dual power hosts in which each power switch has two agents connected to the same power switch. • If this check box is selected, both fencing agents are used concurrently when a host is fenced. This means that both fencing agents have to respond to the Stop command for the host to be stopped; if one agent responds to the Start command, the host will go up. • If this check box is not selected, the fencing agents are used sequentially. This means that to stop or start a host, the primary agent is used first, and if it fails, the secondary agent is used.
Address	The address to access your host's power management device. Either a resolvable hostname or an IP address.
User Name	User account to access the power management device with. You may have to set up a user on the device, or use the default user.
Password	Password for the user accessing the power management device.

Field Name	Description
Type	The type of power management device in your host. Choose one of the following: - **apc** - APC MasterSwitch network power switch. Not for use with APC 5.x power switch devices. - **apc_snmp** - Use with APC 5.x power switch devices. - **bladecenter** - IBM Bladecentre Remote Supervisor Adapter - **cisco_ucs** - Cisco Unified Computing System - **drac5** - Dell Remote Access Controller for Dell computers - **eps** - ePowerSwitch 8M+ network power switch - **ilo**, **ilo2**, **ilo3**, **ilo4** - HP Integrated Lights-Out - **ipmilan** - Intelligent Platform Management Interface and Sun Integrated Lights Out Management devices. - **rsa** - IBM Remote Supervisor Adaptor - **rsb** - Fujitsu-Siemens RSB management interface - **wti** - WTI Network PowerSwitch
Port	The port number used by the power management device to communicate with the host.
Options	Power management device specific options. Give these as 'key=value' or 'key', refer to the documentation of your host's power management device for the options available.
Secure	Tick this check box to allow the power management device to connect securely to the host. This can be done via ssh, ssl, or other authentication protocols depending on and supported by the power management agent.
Source	Specifies whether the host will search within its **cluster** or **data center** for a fencing proxy. Use the **Up** and **Down** buttons to change the sequence in which the resources are used.

Report a bug

9.5.3. SPM Priority Settings Explained

The **SPM** settings table details the information required on the **SPM** tab of the **New Host** or **Edit Host** window.

Table 9.3. SPM settings

Field Name	Description
SPM Priority	Defines the likelihood that the host will be given the role of Storage Pool Manager(SPM). The options are Low, Normal, and High priority, where Low priority means a reduced likelihood of the host being assigned the role of SPM, and High priority increases the likelihood. The default setting is Normal.

Report a bug

9.5.4. Host Console Settings Explained

The **Console** settings table details the information required on the **Console** tab of the **New Host** or **Edit Host** window.

Table 9.4. Console settings

Field Name	Description
Override display address	Select this check box to enable overriding the display addresses of the host. This feature is useful in a case where the hosts are defined by internal IP and are behind a NAT firewall. When a user connects to a virtual machine from outside of the internal network, instead of returning the private address of the host on which the virtual machine is running, a public IP or FQDN (which is resolved in the external network to the public IP) is returned.
Display address	The display address specified here will be used for all virtual machines running on this host. The address must be in the format of a fully qualified domain name or IP.

Report a bug

Part IV. Environment Configuration

Chapter 10. Planning your Data Center

10.1. Workflow Progress — Planning Your Data Center

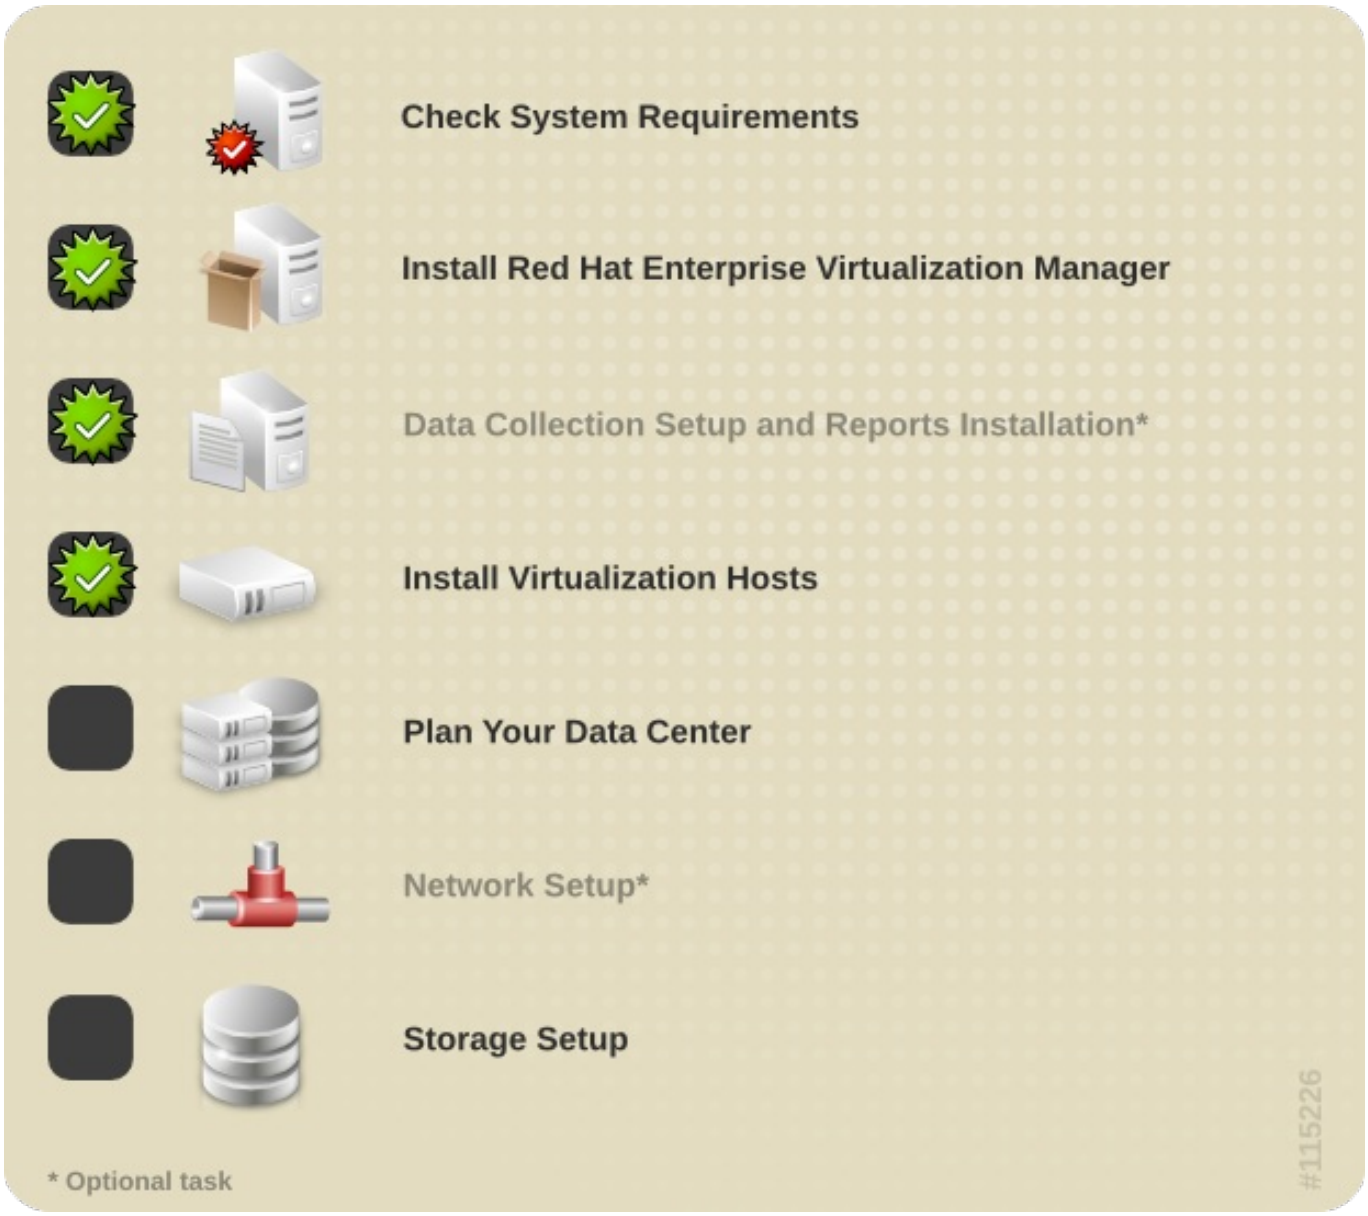

Report a bug

10.2. Planning Your Data Center

Successful planning is essential for a highly available, scalable Red Hat Enterprise Virtualization environment.

Although it is assumed that your solution architect has defined the environment before installation, the following considerations must be made when designing the system.

CPU

Virtual Machines must be distributed across hosts so that enough capacity is available to handle higher than average loads during peak processing. Average target utilization will be 50% of available CPU.

Memory

The Red Hat Enterprise Virtualization page sharing process overcommits up to 150% of physical memory for virtual machines. Therefore, allow for an approximately 30% overcommit.

Networking

When designing the network, it is important to ensure that the volume of traffic produced by storage, remote connections and virtual machines is taken into account. As a general rule, allow approximately 50 MBps per virtual machine.

It is best practice to separate disk I/O traffic from end-user traffic, as this reduces the load on the Ethernet connection and reduces security vulnerabilities by isolating data from the visual stream. For Ethernet networks, it is suggested that bonds (802.3ad) are utilized to aggregate server traffic types.

Note

It is possible to connect both the storage and Hypervisors via a single high performance switch. For this configuration to be effective, the switch must be able to provide 30 GBps on the backplane.

High Availability

The system requires at least two hosts to achieve high availability. This redundancy is useful when performing maintenance or repairs.

Report a bug

10.3. Data Centers

10.3.1. Data Centers in Red Hat Enterprise Virtualization

The data center is the highest level container for all physical and logical resources within a managed virtual environment. The data center is a collection of clusters of Hosts. It owns the logical network (that is, the defined subnets for management, guest network traffic and storage network traffic) and the storage pool.

Red Hat Enterprise Virtualization contains a **`Default`** data center at installation. You can create new data centers that will also be managed from the single Administration Portal. For example, you may choose to have different data centers for different physical locations, business units, or for reasons of security. It is recommended that you do not remove the **`Default`** data center, instead set up new appropriately named data centers.

The system administrator, as the superuser, can manage all aspects of the platform, that is, data centers, storage domains, users, roles and permissions by default; however more specific administrative roles and permissions can be assigned to other users. For example, the enterprise may need a Data Center administrator for a specific data center, or a particular cluster may need an administrator. All system administration roles for physical resources have a hierarchical permission system. For example, a data center administrator will automatically have permission to manage all the objects in that data center - including storage domains, clusters and hosts.

See Also:

- Section 10.3.2, “Creating a New Data Center”

- Section 10.4.1, “Clusters in Red Hat Enterprise Virtualization”

Report a bug

10.3.2. Creating a New Data Center

Summary

This procedure creates a data center in your virtualization environment. The data center requires a functioning cluster, host, and storage domain to operate.

Note

The storage **Type** can be edited until the first storage domain is added to the data center. Once a storage domain has been added, the storage **Type** cannot be changed.

If you set the **Compatibility Version** as **3.1**, it cannot be changed to **3.0** at a later time; version regression is not allowed.

Procedure 10.1. Creating a New Data Center

1. Select the **Data Centers** resource tab to list all data centers in the results list.
2. Click **New** to open the **New Data Center** window.
3. Enter the **Name** and **Description** of the data center.
4. Select the storage **Type**, **Compatibility Version**, and **Quota Mode** of the data center from the drop-down menus.
5. Click **OK** to create the data center and open the **New Data Center - Guide Me** window.
6. The **Guide Me** window lists the entities that need to be configured for the data center. Configure these entities or postpone configuration by clicking the **Configure Later** button; configuration can be resumed by selecting the data center and clicking the **Guide Me** button.

Result

The new data center is added to the virtualization environment. It will remain **Uninitialized** until a cluster, host, and storage domain is configured for it; use **Guide Me** to configure these entities.

See Also:

- Section 10.3.1, “Data Centers in Red Hat Enterprise Virtualization”

Report a bug

10.4. Clusters

10.4.1. Clusters in Red Hat Enterprise Virtualization

A cluster is a collection of physical hosts that share similar characteristics and work together to

provide computing resources in a highly available manner. In Red Hat Enterprise Virtualization the cluster must contain physical hosts that share the same storage domains and have the same type of CPU. Because virtual machines can be migrated across hosts in the same cluster, the cluster is the highest level at which power and load-sharing policies can be defined. The Red Hat Enterprise Virtualization platform contains a **Default** cluster in the **Default** data center at installation time.

Every cluster in the system must belong to a data center, and every host in the system must belong to a cluster. This enables the system to dynamically allocate a virtual machine to any host in the cluster, according to policies defined on the **Cluster** tab, thus maximizing memory and disk space, as well as virtual machine uptime.

At any given time, after a virtual machine runs on a specific host in the cluster, the virtual machine can be migrated to another host in the cluster using **Migrate**. This can be very useful when a host must be shut down for maintenance. The migration to another host in the cluster is transparent to the user, and the user continues working as usual. Note that a virtual machine can not be migrated to a host outside the cluster.

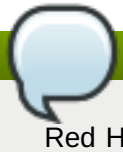

Note

Red Hat Enterprise Virtualization 3.1 supports the use of clusters to manage Gluster storage bricks, in addition to virtualization hosts. To begin managing Gluster storage bricks, create a cluster with the **Enable Gluster Service** option selected. For further information on Gluster storage bricks, refer to the Red Hat Enterprise Storage *Administration Guide*, available at http://docs.redhat.com/.

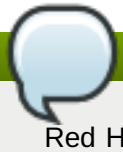

Note

Red Hat Enterprise Virtualization supports **Memory Optimization** by enabling and tuning Kernel Same-page Merging (KSM) on the virtualization hosts in the cluster. For more information on KSM refer to the Red Hat Enterprise Linux 6 *Virtualization Administration Guide*.

See Also:

- Section 10.4.2, “Creating a New Cluster”
- Section 10.3.1, “Data Centers in Red Hat Enterprise Virtualization”

Report a bug

10.4.2. Creating a New Cluster

Summary

A data center can contain multiple clusters, and a cluster can contain multiple hosts. All hosts in a cluster must be of the same CPU type (Intel or AMD). It is recommended that you create your hosts before you create your cluster to ensure CPU type optimization. However, you can configure the hosts at a later time using the **Guide Me** button.

Procedure 10.2. Creating a New Cluster

1. Select the **Clusters** resource tab.

2. Click **New** to open the **New Cluster** window.
3. Select the **Data Center** the cluster will belong to from the drop-down list.
4. Enter the **Name** and **Description** of the cluster.
5. Select the **CPU Name** and **Compatibility Version** from the drop-down lists. It is important to match the CPU processor family with the minimum CPU processor type of the hosts you intend to attach to the cluster, otherwise the host will be non-operational.
6. Select either the **Enable Virt Service** or **Enable Gluster Service** radio box depending on whether the cluster should be populated with virtual machine hosts or Gluster-enabled nodes. Note that you cannot add Red Hat Enterprise Virtualization Hypervisor hosts to a Gluster-enabled cluster.
7. Click the **Optimization** tab to select the memory page sharing threshold for the cluster, and optionally enable CPU thread handling and memory ballooning on the hosts in the cluster.
8. Click the **Cluster Policy** tab to optionally configure a power policy, scheduler optimization settings, and enable trusted service for hosts in the cluster.
9. Click the **Resilience Policy** tab to select the virtual machine migration policy.
10. Click **OK** to create the cluster and open the **New Cluster - Guide Me** window.
11. The **Guide Me** window lists the entities that need to be configured for the cluster. Configure these entities or postpone configuration by clicking the **Configure Later** button; configuration can be resumed by selecting the cluster and clicking the **Guide Me** button.

Result

The new cluster is added to the virtualization environment.

See Also:

- Section 10.4.1, “Clusters in Red Hat Enterprise Virtualization”

Report a bug

10.4.3. Enabling Gluster Processes on Red Hat Storage Nodes

Summary

This procedure explains how to allow Gluster processes on Red Hat Storage Nodes.

1. In the Navigation Pane, select the **Clusters** tab.
2. Select **New**.
3. Select the "Enable Gluster Service" radio button. Provide the address, SSH fingerprint, and password as necessary. The address and password fields can be filled in only when the **Import existing Gluster configuration** check box is selected.

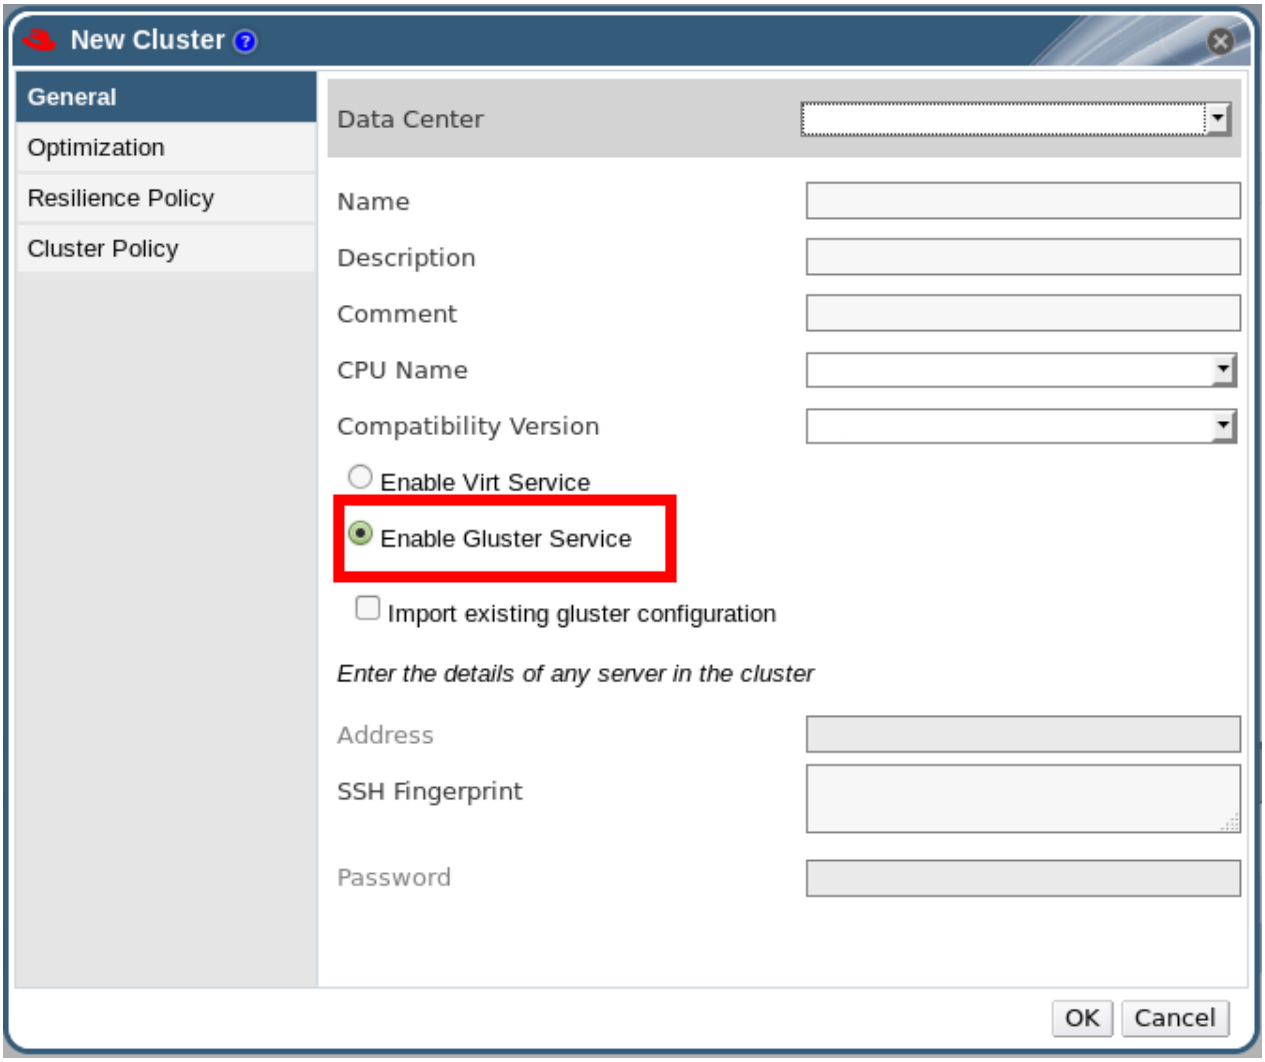

Figure 10.1. Selecting the "Enable Gluster Service" Radio Button

4. Click **OK**.

Result

It is now possible to add Red Hat Storage nodes to the Gluster cluster, and to mount Gluster volumes as storage domains. **iptables** rules no longer block storage domains from being added to the cluster.

Report a bug

Chapter 11. Network Setup

11.1. Workflow Progress — Network Setup

Report a bug

11.2. Networking in Red Hat Enterprise Virtualization

Red Hat Enterprise Virtualization uses networking to support almost every aspect of operations. Storage, host management, user connections, and virtual machine connectivity, for example, all rely on a well planned and configured network to deliver optimal performance. Setting up networking is a vital prerequisite for a Red Hat Enterprise Virtualization environment because it is much simpler to plan for your projected networking requirements and implement your network accordingly than it is to discover your networking requirements through use and attempt to alter your network configuration retroactively.

It is however possible to deploy a Red Hat Enterprise Virtualization environment with no consideration given to networking at all. Simply ensuring that each physical machine in the environment has at least one *Network Interface Controller* (NIC) is enough to begin using Red Hat Enterprise Virtualization. While it is true that this approach to networking will provide a functional

environment, it will not provide an optimal environment. As network usage varies by task or action, grouping related tasks or functions into specialized networks can improve performance while simplifying the troubleshooting of network issues.

Red Hat Enterprise Virtualization separates network traffic by defining logical networks. Logical networks define the path that a selected network traffic type must take through the network. They are created to isolate network traffic by functionality or virtualize a physical topology.

The **rhevm** logical network is created by default and labeled as the **Management**. The **rhevm** logical network is intended for management traffic between the Red Hat Enterprise Virtualization Manager and virtualization hosts. You are able to define additional logical networks to segregate:

- Display related network traffic.
- General virtual machine network traffic.
- Storage related network traffic.

For optimal performance it is recommended that these traffic types be separated using logical networks. Logical networks may be supported using physical devices such as NICs or logical devices, such as network bonds. It is not necessary to have one device for each logical network as multiple logical networks are able to share a single device. This is accomplished using Virtual LAN (VLAN) tagging to isolate network traffic. To make use of this facility VLAN tagging must also be supported at the switch level.

The limits that apply to the number of logical networks that you may define in a Red Hat Enterprise Virtualization environment are:

- The number of logical networks attached to a host is limited to the number of available network devices combined with the maximum number of Virtual LANs (VLANs) which is 4096.
- The number of logical networks in a cluster is limited to the number of logical networks that can be attached to a host as networking must be the same for all hosts in a cluster.
- The number of logical networks in a data center is limited only by the number of clusters it contains in combination with the number of logical networks permitted per cluster.

Note

From Red Hat Enterprise Virtualization 3.3, network traffic for migrating virtual machines has been separated from network traffic for communication between the Manager and hosts. This prevents hosts from becoming non-responsive when importing or migrating virtual machines.

Note

A familiarity with the network concepts and their use is highly recommended when planning and setting up networking in a Red Hat Enterprise Virtualization environment. This document does not describe the concepts, protocols, requirements or general usage of networking. It is recommended that you read your network hardware vendor's guides for more information on managing networking.

Important

Additional care must be taken when modifying the properties of the **rhevm** network. Incorrect changes to the properties of the **rhevm** network may cause hosts to become temporarily unreachable.

Important

If you plan to use Red Hat Enterprise Virtualization nodes to provide any services, remember that the services will stop if the Red Hat Enterprise Virtualization environment stops operating.

This applies to all services, but you should be fully aware of the hazards of running the following on Red Hat Enterprise Virtualization:

- Directory Services
- DNS
- Storage

See Also:

- Section 11.3, “Logical Networks”
- Section 11.5, “Bonds”

Report a bug

11.3. Logical Networks

11.3.1. Creating a New Logical Network in a Data Center or Cluster

Summary

Create a logical network and define its use in a data center, or in clusters in a data center.

Procedure 11.1. Creating a New Logical Network in a Data Center or Cluster

1. Use the **Data Centers** or **Clusters** resource tabs, tree mode, or the search function to find and select a data center or cluster in the results list.
2. Click the **Logical Networks** tab of the details pane to list the existing logical networks.
3. From the **Data Centers** details pane, click **New** to open the **New Logical Network** window.

   From the **Clusters** details pane, click **Add Network** to open the **New Logical Network** window.
4. Enter a **Name**, **Description** and **Comment** for the logical network.
5. In the **Export** section, select the **Create on external provider** check box to create the logical network on an external provider. Select the external provider from the **External Provider** drop-down list and enter a **Network Label** for the logical network.

6. Select the **Enable VLAN tagging**, **VM network** and **Override MTU** to enable these options.
7. From the **Cluster** tab, select the clusters to which the network will be assigned. You can also specify whether the logical network will be a required network.
8. From the **Profiles** tab, add vNIC profiles to the logical network as required.
9. Click **OK**.

Result

You have defined this logical network as a resource required by a cluster or clusters in the data center. You can now add this resource to the hosts in the cluster.

Note

When creating a new logical network or making changes to an existing logical network that is used as a display network, any running virtual machines that use that network must be rebooted before the network becomes available or the changes are applied.

See Also:

- Section 10.3.1, “Data Centers in Red Hat Enterprise Virtualization”
- Section 10.4.1, “Clusters in Red Hat Enterprise Virtualization”

Report a bug

11.3.2. Editing Host Network Interfaces and Adding Logical Networks to Hosts

Summary

You can change the settings of host network interfaces. Moving the **rhevm** management logical network between interfaces, and adding a newly created logical network to a network interface are common reasons to edit host networking.

Procedure 11.2. Editing Host Network Interfaces and Adding Logical Networks to Hosts

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results.
2. Click the **Network Interfaces** tab in the details pane to list the network interfaces attached to the host and their configurations.
3. Click the **Setup Host Networks** button to open the **Setup Host Networks** window.

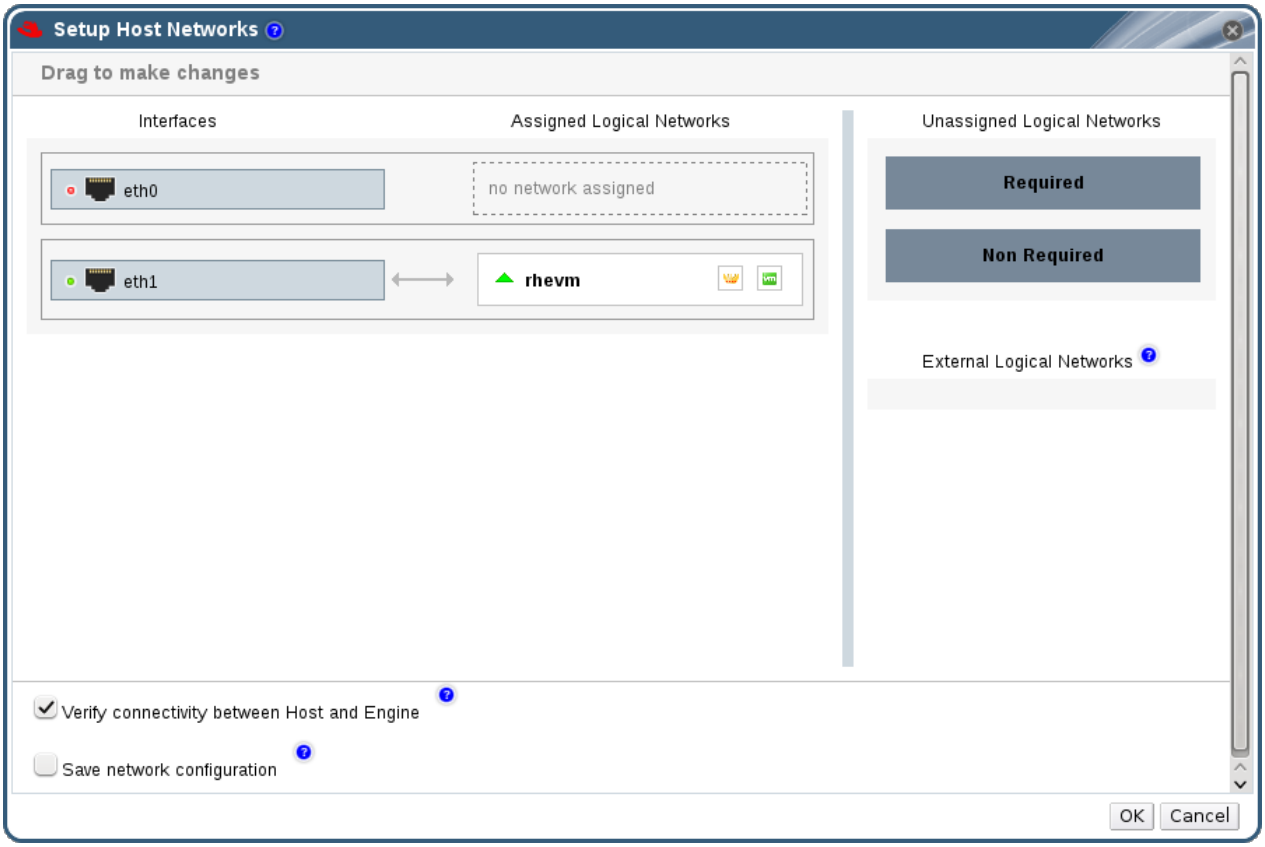

Figure 11.1. The Setup Host Networks window

4. Attach a logical network to a network interface by selecting and dragging a logical network into the **Assigned Logical Networks** area next to the network interface.

   Alternatively, right-click the logical network and select a network interface from the drop-down menu.

5. Edit the logical networks by hovering your cursor over an assigned logical network and clicking the pencil icon to open the **Edit Management Network** window.

   If your logical network definition is not synchronized with the network configuration on the host, select the **Sync network** check box.

   Select a **Boot Protocol** from:

   - **None**,
   - **DHCP**, or
   - **Static**.

     If you have chosen **Static**, provide the **IP**, **Subnet Mask**, and the **Gateway**.

   Click **OK**.

6. Select the **Verify connectivity between Host and Engine** check box to run a network check; this will only work if the host is in maintenance mode.

7. Select the **Save network configuration** check box if you want these network changes to be made persistent when the environment is rebooted.

8. Click **OK** to implement the changes and close the window.

Result

You have assigned logical networks to network interfaces and configured the host network.

Note

If not all network interface cards for the host are displayed, click the **Refresh Capabilities** button to update the list of network interface cards available for that host.

Report a bug

11.3.3. Explanation of Settings and Controls in the General Tab of the New Logical Network and Edit Logical Network Windows

The table below describes the settings for the **General** tab of the **New Logical Network** and **Edit Logical Network** window.

Table 11.1. New Logical Network and Edit Logical Network Settings

Field Name	Description
Name	The name of the logical network. This text field has a 15-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
Description	The description of the logical network. This field is recommended but not mandatory.
Comment	A field for adding plain text, human-readable comments regarding the logical network.
Export	Allows you to export the logical network to an OpenStack Network Service that has been added to the Manager as an external provider. **External Provider** - Allows you to select the external provider on which the logical network will be created. **Network Label** - Allows you to specify the label of the logical network, such as **eth0**.
Enable VLAN tagging	VLAN tagging is a security feature that gives all network traffic carried on the logical network a special characteristic. VLAN-tagged traffic cannot be read by interfaces that do not also have that characteristic. Use of VLANs on logical networks also allows a single network interface to be associated with multiple, differently VLAN-tagged logical networks. Enter a numeric value in the text entry field if VLAN tagging is enabled.

Field Name	Description
VM Network	Select this option if only virtual machines use this network. If the network is used for traffic that does not involve virtual machines, such as storage communications, do not select this check box.
Override MTU	Set a custom maximum transmission unit for the logical network. You can use this to match the MTU supported by your new logical network to the MTU supported by the hardware it interfaces with. Enter a numeric value in the text entry field if MTU override is enabled.

Report a bug

11.3.4. Editing a Logical Network

Summary

Edit the settings of a logical network.

Procedure 11.3. Editing a Logical Network

1. Use the **Data Centers** resource tab, tree mode, or the search function to find and select the data center of the logical network in the results list.
2. Click the **Logical Networks** tab in the details pane to list the logical networks in the data center.
3. Select a logical network and click **Edit** to open the **Edit Logical Network** window.
4. Edit the necessary settings.
5. Click **OK** to save the changes.

Result

You have updated the settings of your logical network.

See Also:

- Section 11.3.3, “Explanation of Settings and Controls in the General Tab of the New Logical Network and Edit Logical Network Windows”

Report a bug

11.3.5. Designate a Specific Traffic Type for a Logical Network with the Manage Networks Window

Prerequisites:

- Section 11.3.1, “Creating a New Logical Network in a Data Center or Cluster”

Summary

Specify the traffic type for the logical network to optimize the network traffic flow.

Procedure 11.4. Assigning or Unassigning a Logical Network to a Cluster

1. Use the **Clusters** resource tab, tree mode, or the search function to find and select the cluster in the results list.
2. Select the **Logical Networks** tab in the details pane to list the logical networks assigned to the cluster.
3. Click **Manage Networks** to open the **Manage Networks** window.

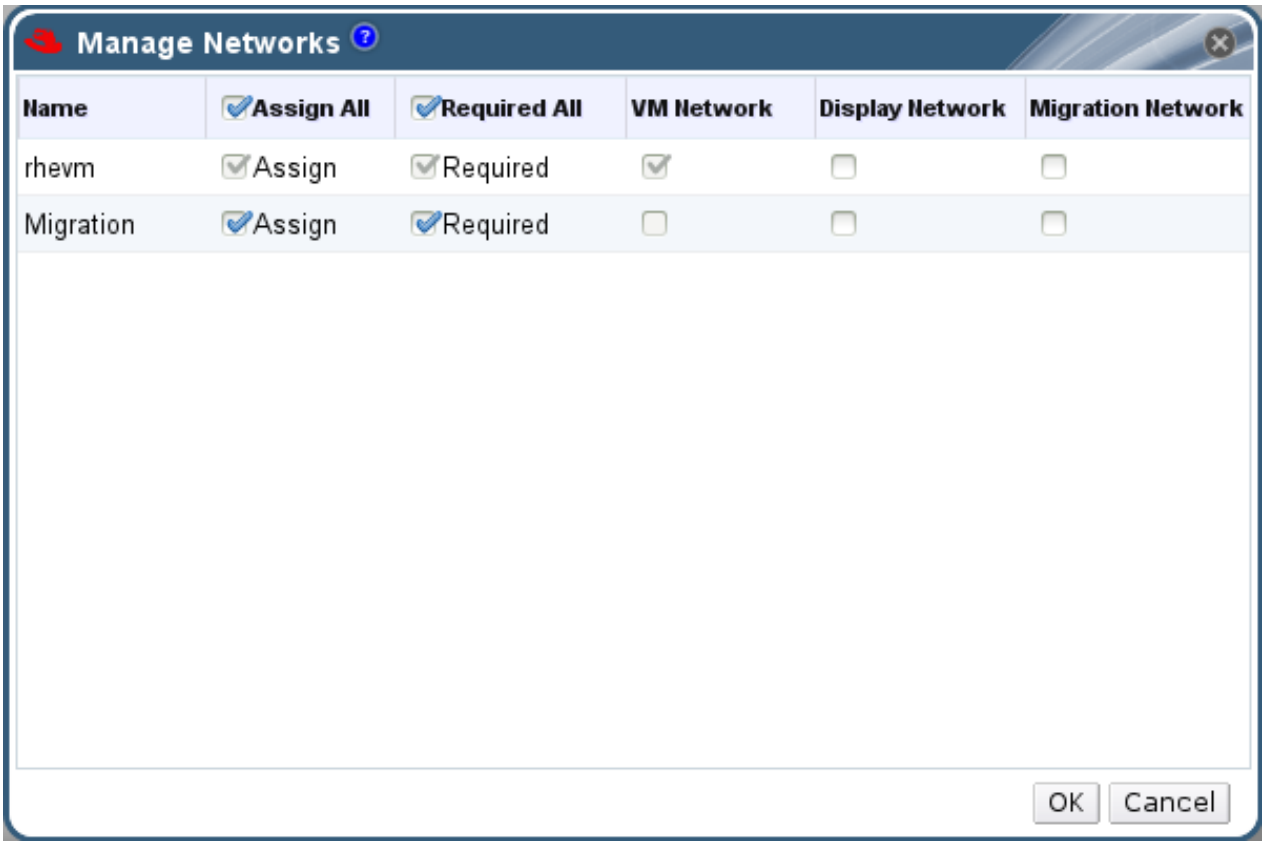

Figure 11.2. Manage Networks

4. Select appropriate check boxes.
5. Click **OK** to save the changes and close the window.

Result

You have optimized the network traffic flow by assigning a specific type of traffic to be carried on a specific logical network.

Note

Networks offered by external providers cannot be used as display networks.

See Also:

- Section 11.3.2, “Editing Host Network Interfaces and Adding Logical Networks to Hosts”

- Section 11.3.6, “Explanation of Settings in the Manage Networks Window”

Report a bug

11.3.6. Explanation of Settings in the Manage Networks Window

The table below describes the settings for the **Manage Networks** window.

Table 11.2. Manage Networks Settings

Field	Description/Action
Assign	Assigns the logical network to all hosts in the cluster.
Required	A logical network becomes operational when it is attached to an active NIC on all hosts in the cluster.
VM Network	The logical network carries the virtual machine network traffic.
Display Network	The logical network carries the virtual machine SPICE and virtual network controller traffic.
Migration Network	The logical network carries virtual machine and storage migration traffic.

Report a bug

11.3.7. Adding Multiple VLANs to a Single Network Interface Using Logical Networks

Prerequisites:

- Section 11.3.1, “Creating a New Logical Network in a Data Center or Cluster”

Summary

Multiple VLANs can be added to a single network interface to separate traffic on the one host.

Important

You must have created more than one logical network, all with the **Enable VLAN tagging** check box selected in the **New Logical Network** or **Edit Logical Network** windows.

Procedure 11.5. Adding Multiple VLANs to a Network Interface using Logical Networks

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select in the results list a host associated with the cluster to which your VLAN-tagged logical networks are assigned.
2. Click the **Network Interfaces** tab in the details pane to list the physical network interfaces attached to the data center.
3. Click **Setup Host Networks** to open the **Setup Host Networks** window.

4. Drag your VLAN-tagged logical networks into the **Assigned Logical Networks** area next to the physical network interface. The physical network interface can have multiple logical networks assigned due to the VLAN tagging.

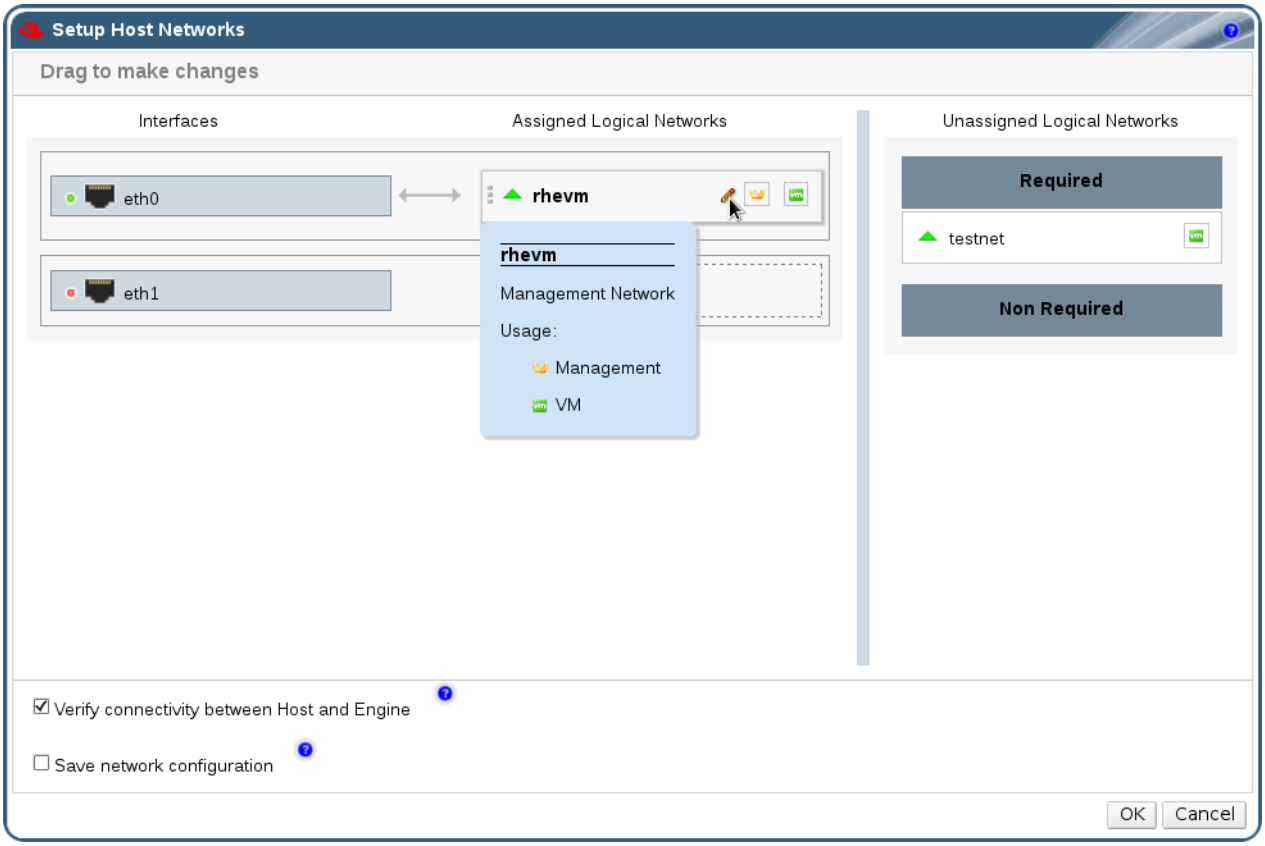

Figure 11.3. Setup Host Networks

5. Edit the logical networks by hovering your cursor over an assigned logical network and clicking the pencil icon to open the **Edit Network** window.

   If your logical network definition is not synchronized with the network configuration on the host, select the **Sync network** check box.

   Select a **Boot Protocol** from:

   - **None**,
   - **DHCP**, or
   - **Static**,

     Provide the **IP** and **Subnet Mask**.

   Click **OK**.

6. Select the **Verify connectivity between Host and Engine** check box to run a network check; this will only work if the host is in maintenance mode.

7. Select the **Save network configuration** check box

8. Click **OK**.

Add the logical network to each host in the cluster by editing a NIC on each host in the cluster. After this is done, the network will become operational

Result

You have added multiple VLAN-tagged logical networks to a single interface. This process can be repeated multiple times, selecting and editing the same network interface each time on each host to add logical networks with different VLAN tags to a single network interface.

Report a bug

11.3.8. Multiple Gateways

Summary

Users can define the gateway, along with the IP address and subnet mask, for a logical network. This is necessary when multiple networks exist on a host and traffic should be routed through the specified network, rather than the default gateway.

If multiple networks exist on a host and the gateways are not defined, return traffic will be routed through the default gateway, which may not reach the intended destination. This would result in users being unable to ping the host.

Red Hat Enterprise Virtualization 3.3 handles multiple gateways automatically whenever an interface goes up or down.

Procedure 11.6. Viewing or Editing the Gateway for a Logical Network

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click the **Network Interfaces** tab in the details pane to list the network interfaces attached to the host and their configurations.
3. Click the **Setup Host Networks** button to open the **Setup Host Networks** window.
4. Hover your cursor over an assigned logical network and click the pencil icon to open the **Edit Management Network** window.

Result

The **Edit Management Network** window displays the network name, the boot protocol, and the IP, subnet mask, and gateway addresses. The address information can be manually edited by selecting a **Static** boot protocol.

Report a bug

11.4. Using the Networks Tab

The **Networks** resource tab provides a central location for users to perform network-related operations and search for networks based on each network's property or association with other resources.

All networks in the Red Hat Enterprise Virtualization environment display in the results list of the **Networks** tab. The **New**, **Edit** and **Remove** buttons allow you to create, change the properties of, and delete logical networks within data centers.

Click on each network name and use the **Clusters**, **Hosts**, **Virtual Machines**, **Templates**, and **Permissions** tabs in the details pane to perform functions including:

- Attaching or detaching the networks to clusters and hosts
- Removing network interfaces from virtual machines and templates
- Adding and removing permissions for users to access and manage networks

These functions are also accessible through each individual resource tab.

Report a bug

11.4.1. Importing Networks from External Providers

Summary

If an external provider offering networking services has been registered in the Manager, the networks provided by that provider can be imported into the Manager and used by virtual machines.

Procedure 11.7. Importing a Network

1. Click on the **Networks** tab.
2. Click the **Import** button. The **Import Networks** window appears.
3. From the **Network Provider** drop-down list, select a provider. The networks offered by that provider are automatically discovered and display in the **Provider Networks** list.
4. Select the network to import in the **Provider Networks** list and click the down arrow to move the network into the **Networks to Import** list.
5. Click the **Import** button.

Result

The selected networks are imported and can now be used within the Manager.

Important

External provider discovery and importing are Technology Preview features. Technology Preview features are not fully supported under Red Hat Subscription Service Level Agreements (SLAs), may not be functionally complete, and are not intended for production use. However, these features provide early access to upcoming product innovations, enabling customers to test functionality and provide feedback during the development process.

Report a bug

11.4.2. Limitations to Importing Networks from External Providers

While networks offered by external providers can be imported into the Manager, the following limitations apply to their usage:

- Networks offered by external providers must be used as virtual machine networks.
- Networks offered by external providers cannot be used as display networks.
- The same network can be imported more than once, but only to different data centers.

- Networks offered by external providers cannot be edited in the Manager. This is because the management of such networks is the responsibility of the external providers.
- Port mirroring is not available for virtual NIC connected to networks offered by external providers.
- If a virtual machine uses a network offered by an external provider, that provider cannot be deleted from the Manager while the network is still in use by the virtual machine.
- Networks offered by external providers are non-required. As such, scheduling for clusters in which such networks have been imported will not take those networks into account during host selection. Moreover, it is the responsibility of the user to ensure the availability of the network on hosts in clusters in which such networks have been imported.

Important

External provider discovery and importing are Technology Preview features. Technology Preview features are not fully supported under Red Hat Subscription Service Level Agreements (SLAs), may not be functionally complete, and are not intended for production use. However, these features provide early access to upcoming product innovations, enabling customers to test functionality and provide feedback during the development process.

Report a bug

11.5. Bonds

11.5.1. Bonding Logic in Red Hat Enterprise Virtualization

The Red Hat Enterprise Virtualization Manager Administration Portal allows you to create bond devices using a graphical interface. There are several distinct bond creation scenarios, each with its own logic.

Two factors that affect bonding logic are:

- Are either of the devices already carrying logical networks?
- Are the devices carrying compatible logical networks? A single device cannot carry both VLAN tagged and non-VLAN tagged logical networks.

Table 11.3. Bonding Scenarios and Their Results

Bonding Scenario	Result
NIC + NIC	The **Create New Bond** window is displayed, and you can configure a new bond device. If the network interfaces carry incompatible logical networks, the bonding operation fails until you detach incompatible logical networks from the devices forming your new bond.

Bonding Scenario	Result
NIC + Bond	The NIC is added to the bond device. Logical networks carried by the NIC and the bond are all added to the resultant bond device if they are compatible. If the bond devices carry incompatible logical networks, the bonding operation fails until you detach incompatible logical networks from the devices forming your new bond.
Bond + Bond	If the bond devices are not attached to logical networks, or are attached to compatible logical networks, a new bond device is created. It contains all of the network interfaces, and carries all logical networks, of the component bond devices. The **Create New Bond** window is displayed, allowing you to configure your new bond. If the bond devices carry incompatible logical networks, the bonding operation fails until you detach incompatible logical networks from the devices forming your new bond.

Report a bug

11.5.2. Bonding Modes

Red Hat Enterprise Virtualization supports the following common bonding modes:

- Mode 1 (active-backup policy) sets all interfaces to the backup state while one remains active. Upon failure on the active interface, a backup interface replaces it as the only active interface in the bond. The MAC address of the bond in mode 1 is visible on only one port (the network adapter), to prevent confusion for the switch. Mode 1 provides fault tolerance and is supported in Red Hat Enterprise Virtualization.
- Mode 2 (XOR policy) selects an interface to transmit packages to based on the result of an XOR operation on the source and destination MAC addresses modulo NIC slave count. This calculation ensures that the same interface is selected for each destination MAC address used. Mode 2 provides fault tolerance and load balancing and is supported in Red Hat Enterprise Virtualization.
- Mode 4 (IEEE 802.3ad policy) creates aggregation groups for which included interfaces share the speed and duplex settings. Mode 4 uses all interfaces in the active aggregation group in accordance with the IEEE 802.3ad specification and is supported in Red Hat Enterprise Virtualization.
- Mode 5 (adaptive transmit load balancing policy) ensures the outgoing traffic distribution is according to the load on each interface and that the current interface receives all incoming traffic. If the interface assigned to receive traffic fails, another interface is assigned the receiving role instead. Mode 5 is supported in Red Hat Enterprise Virtualization.

Report a bug

11.5.3. Creating a Bond Device Using the Administration Portal

Summary

You can bond compatible network devices together. This type of configuration can increase available bandwidth and reliability. You can bond multiple network interfaces, pre-existing bond devices, and combinations of the two.

A bond cannot carry both vlan tagged and non-vlan traffic.

Procedure 11.8. Creating a Bond Device using the Administration Portal

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click the **Network Interfaces** tab in the details pane to list the physical network interfaces attached to the host.
3. Click **Setup Host Networks** to open the **Setup Host Networks** window.
4. Select and drag one of the devices over the top of another device and drop it to open the **Create New Bond** window. Alternatively, right-click the device and select another device from the drop-down menu.

   If the devices are incompatible, for example one is vlan tagged and the other is not, the bond operation fails with a suggestion on how to correct the compatibility issue.

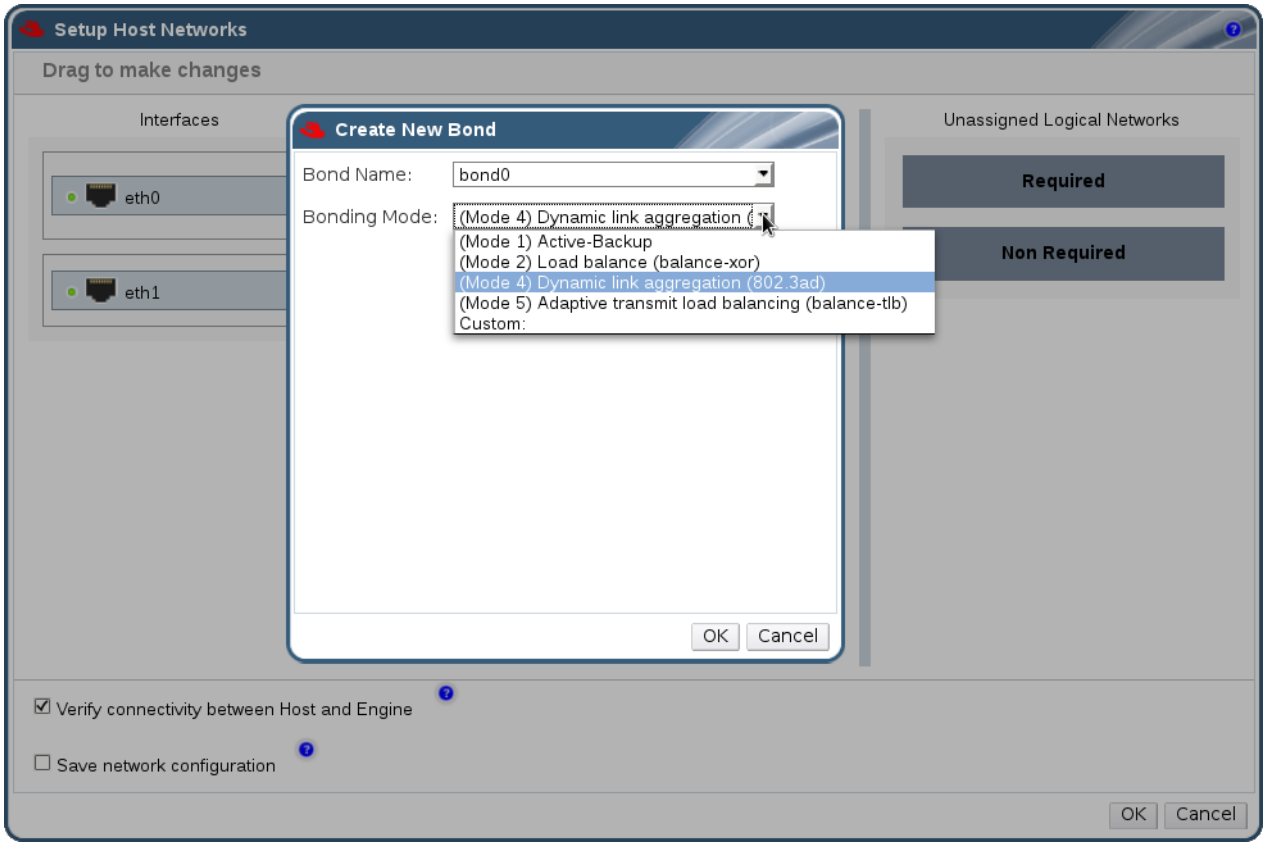

Figure 11.4. Bond Devices Window

5. Select the **Bond Name** and **Bonding Mode** from the drop-down menus.

   Bonding modes 1, 2, 4, and 5 can be selected. Any other mode can be configured using the **Custom** option.

6. Click **OK** to create the bond and close the **Create New Bond** window.
7. Assign a logical network to the newly created bond device.
8. Optionally choose to **Verify connectivity between Host and Engine** and **Save network configuration**.
9. Click **OK** accept the changes and close the **Setup Host Networks** window.

Result:

Your network devices are linked into a bond device and can be edited as a single interface. The bond device is listed in the **Network Interfaces** tab of the details pane for the selected host.

Bonding must be enabled for the ports of the switch used by the host. The process by which bonding is enabled is slightly different for each switch; consult the manual provided by your switch vendor for detailed information on how to enable bonding.

Report a bug

11.5.4. Example Uses of Custom Bonding Options with Host Interfaces

You can create customized bond devices by selecting **Custom** from the **Bonding Mode** of the **Create New Bond** window. The following examples should be adapted for your needs. For a comprehensive list of bonding options and their descriptions, see the *Linux Ethernet Bonding Driver HOWTO* on Kernel.org.

Example 11.1. xmit_hash_policy

This option defines the transmit load balancing policy for bonding modes 2 and 4. For example, if the majority of your traffic is between many different IP addresses, you may want to set a policy to balance by IP address. You can set this load-balancing policy by selecting a **Custom** bonding mode, and entering the following into the text field:

```
mode=4, xmit_hash_policy=layer2+3
```

Example 11.2. ARP Monitoring

ARP monitor is useful for systems which can't or don't report link-state properly via ethtool. Set an ***arp_interval*** on the bond device of the host by selecting a **Custom** bonding mode, and entering the following into the text field:

```
mode=1, arp_interval=1, arp_ip_target=192.168.0.2
```

Example 11.3. Primary

You may want to designate a NIC with higher throughput as the primary interface in a bond device. Designate which NIC is primary by selecting a **Custom** bonding mode, and entering the following into the text field:

```
mode=1, primary=eth0
```

Report a bug

Chapter 12. Storage Setup

12.1. Workflow Progress — Storage Setup

Report a bug

12.2. Introduction to Storage in Red Hat Enterprise Virtualization

Red Hat Enterprise Virtualization uses a centralized storage system for virtual machine disk images, ISO files and snapshots. Storage networking can be implemented using:

- Network File System (NFS)
- GlusterFS exports
- Other POSIX compliant file systems
- Internet Small Computer System Interface (iSCSI)
- Local storage attached directly to the virtualization hosts

- Fibre Channel Protocol (FCP)
- Parallel NFS (pNFS)

Setting up storage is a prerequisite for a new data center because a data center cannot be initialized unless storage domains are attached and activated.

As a Red Hat Enterprise Virtualization system administrator, you need to create, configure, attach and maintain storage for the virtualized enterprise. You should be familiar with the storage types and their use. Read your storage array vendor's guides, and refer to the *Red Hat Enterprise Linux Storage Administration Guide* for more information on the concepts, protocols, requirements or general usage of storage.

The Red Hat Enterprise Virtualization platform enables you to assign and manage storage using the Administration Portal's **Storage** tab. The **Storage** results list displays all the storage domains, and the details pane shows general information about the domain.

Red Hat Enterprise Virtualization platform has three types of storage domains:

- **Data Domain:** A data domain holds the virtual hard disks and OVF files of all the virtual machines and templates in a data center. In addition, snapshots of the virtual machines are also stored in the data domain.

  The data domain cannot be shared across data centers, and the data domain must be of the same type as the data center. For example, a data center of a iSCSI type, must have an iSCSI data domain.

  You must attach a data domain to a data center before you can attach domains of other types to it.

- **ISO Domain:** ISO domains store ISO files (or logical CDs) used to install and boot operating systems and applications for the virtual machines. An ISO domain removes the data center's need for physical media. An ISO domain can be shared across different data centers.

- **Export Domain:** Export domains are temporary storage repositories that are used to copy and move images between data centers and Red Hat Enterprise Virtualization environments. Export domains can be used to backup virtual machines. An export domain can be moved between data centers, however, it can only be active in one data center at a time.

Important

Support for export storage domains backed by storage on anything other than NFS is being deprecated. While existing export storage domains imported from Red Hat Enterprise Virtualization 2.2 environments remain supported new export storage domains must be created on NFS storage.

Only commence configuring and attaching storage for your Red Hat Enterprise Virtualization environment once you have determined the storage needs of your data center(s).

Important

To add storage domains you must be able to successfully access the Administration Portal, and there must be at least one host connected with a status of **Up**.

See Also:

- Section 12.3.1, “Adding NFS Storage”
- Section 12.3.3, “Adding iSCSI Storage”
- Section 12.3.4, “Adding FCP Storage”
- Section 12.3.5, “Adding Local Storage”
- Section 12.3.6, “Adding POSIX Compliant File System Storage”
- Section 12.4, “Populating the ISO Domain”

Report a bug

12.3. Adding Storage to the Environment

12.3.1. Adding NFS Storage

12.3.1.1. Preparing NFS Storage

Summary

These steps must be taken to prepare an NFS file share on a server running Red Hat Enterprise Linux 6 for use with Red Hat Enterprise Virtualization.

Procedure 12.1. Preparing NFS Storage

1. **Install *nfs-utils***

   NFS functionality is provided by the *nfs-utils* package. Before file shares can be created, check that the package is installed by querying the RPM database for the system:

   ```
   $ rpm -qi nfs-utils
   ```

   If the *nfs-utils* package is installed then the package information will be displayed. If no output is displayed then the package is not currently installed. Install it using **yum** while logged in as the **root** user:

   ```
   # yum install nfs-utils
   ```

2. **Configure Boot Scripts**

   To ensure that NFS shares are always available when the system is operational both the **nfs** and **rpcbind** services must start at boot time. Use the **chkconfig** command while logged in as **root** to modify the boot scripts.

   ```
   # chkconfig --add rpcbind
   # chkconfig --add nfs
   # chkconfig rpcbind on
   # chkconfig nfs on
   ```

   Once the boot script configuration has been done, start the services for the first time.

```
# service rpcbind start
# service nfs start
```

3. **Create Directory**

   Create the directory you wish to share using NFS.

   ```
   # mkdir /exports/iso
   ```

   Replace */exports/iso* with the name, and path of the directory you wish to use.

4. **Export Directory**

   To be accessible over the network using NFS the directory must be exported. NFS exports are controlled using the **/etc/exports** configuration file. Each export path appears on a separate line followed by a tab character and any additional NFS options. Exports to be attached to the Red Hat Enterprise Virtualization Manager must have the read, and write, options set.

   To grant read, and write access to **/exports/iso** using NFS for example you add the following line to the **/etc/exports** file.

   ```
   /exports/iso        *(rw)
   ```

   Again, replace */exports/iso* with the name, and path of the directory you wish to use.

5. **Reload NFS Configuration**

   For the changes to the **/etc/exports** file to take effect the service must be told to reload the configuration. To force the service to reload the configuration run the following command as **root**:

   ```
   # service nfs reload
   ```

6. **Set Permissions**

   The NFS export directory must be configured for read write access and must be owned by vdsm:kvm. If these users do not exist on your external NFS server use the following command, assuming that */exports/iso* is the directory to be used as an NFS share.

   ```
   # chown -R 36:36 /exports/iso
   ```

   The permissions on the directory must be set to allow read and write access to both the owner and the group. The owner should also have execute access to the directory. The permissions are set using the **chmod** command. The following command arguments set the required permissions on the */exports/iso* directory.

   ```
   # chmod 0755 /exports/iso
   ```

Result

The NFS file share has been created, and is ready to be attached by the Red Hat Enterprise Virtualization Manager.

Report a bug

12.3.1.2. Attaching NFS Storage

Prerequisites:

- Section 12.3.1.1, “Preparing NFS Storage”

Summary

An NFS type **Storage Domain** is a mounted NFS share that is attached to a data center. It is used to provide storage for virtualized guest images and ISO boot media. Once NFS storage has been exported it must be attached to the Red Hat Enterprise Virtualization Manager using the Administration Portal.

NFS data domains can be added to NFS data centers. You can add NFS, ISO, and export storage domains to data centers of any type.

Procedure 12.2. Attaching NFS Storage

1. Click the **Storage** resource tab to list the existing storage domains.
2. Click **New Domain** to open the **New Domain** window.

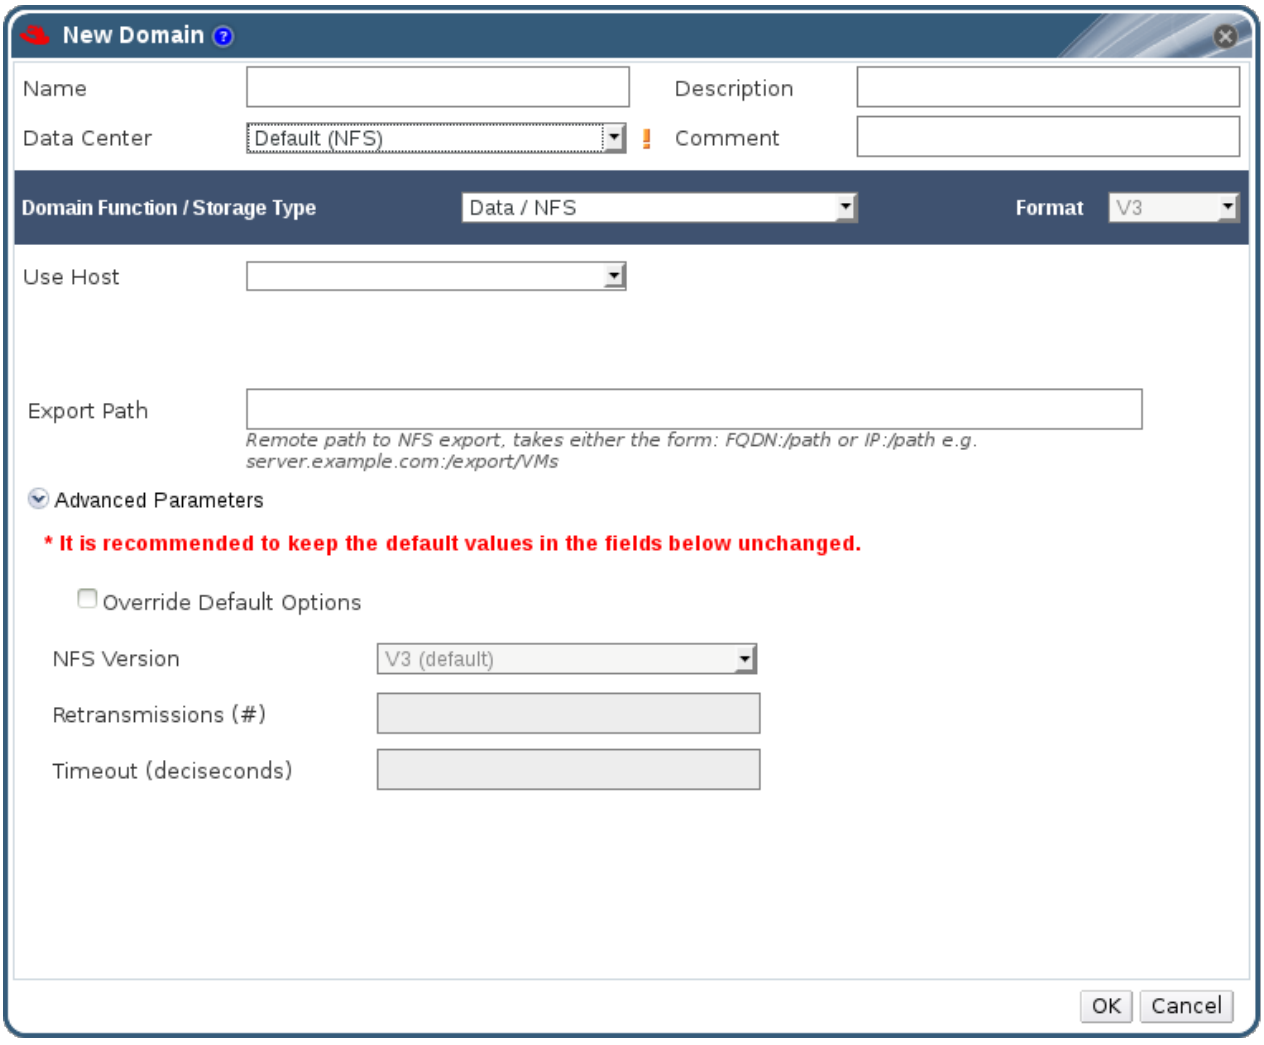

Figure 12.1. NFS Storage

3. Enter the **Name** of the storage domain.
4. Select the **Data Center**, **Domain Function / Storage Type**, and **Use Host** from the drop-down menus.

   If applicable, select the **Format** from the drop-down menu.
5. Enter the **Export Path** to be used for the storage domain.

   The export path should be in the format of **192.168.0.10:/data or domain.example.com:/data**
6. Click **Advanced Parameters** to enable further configurable settings. It is recommended that the values of these parameters not be modified.

Important

All communication to the storage domain is from the selected host and not directly from the Red Hat Enterprise Virtualization Manager. At least one active host must be attached to the chosen Data Center before the storage is configured.

7. Click **OK** to create the storage domain and close the window.

Result

The new NFS data domain is displayed on the **Storage** tab with a status of **Locked** while the disk prepares. It is automatically attached to the data center upon completion.

Report a bug

12.3.2. Adding pNFS Storage

12.3.2.1. Preparing pNFS Storage

Support for Parallel NFS (pNFS) as part of the NFS v4.1 standard is available as of Red Hat Enterprise Linux 6.4. The pNFS architecture improves the scalability of NFS, with possible improvements to performance. That is, when a server implements pNFS as well, a client is able to access data through multiple servers concurrently. The pNFS protocol supports three storage protocols or layouts: files, objects, and blocks. Red Hat Enterprise Linux 6.4 supports only the "files" layout type.

To enable support for pNFS functionality, use one of the following mount options on mounts from a pNFS-enabled server:

```
-o minorversion=1
```

or

```
-o v4.1
```

Set the permissions of the pNFS path so that Red Hat Enterprise Virtualization can access them:

```
# chown 36:36 [path to pNFS resource]
```

After the server is pNFS-enabled, the nfs_layout_nfsv41_files kernel is automatically loaded on the first mount. Verify that the module was loaded:

```
$ lsmod | grep nfs_layout_nfsv41_files
```

Another way to verify a successful NFSv4.1 mount is with the mount command. The mount entry in the output should contain minorversion=1.

Report a bug

12.3.2.2. Attaching pNFS Storage

Summary

A pNFS type **Storage Domain** is a mounted pNFS share attached to a data center. It provides storage for virtualized guest images and ISO boot media. After you have exported pNFS storage, it must be attached to the Red Hat Enterprise Virtualization Manager using the Administration Portal.

Procedure 12.3. Attaching pNFS Storage

1. Click the **Storage** resource tab to list the existing storage domains.
2. Click **New Domain** to open the **New Domain** window.

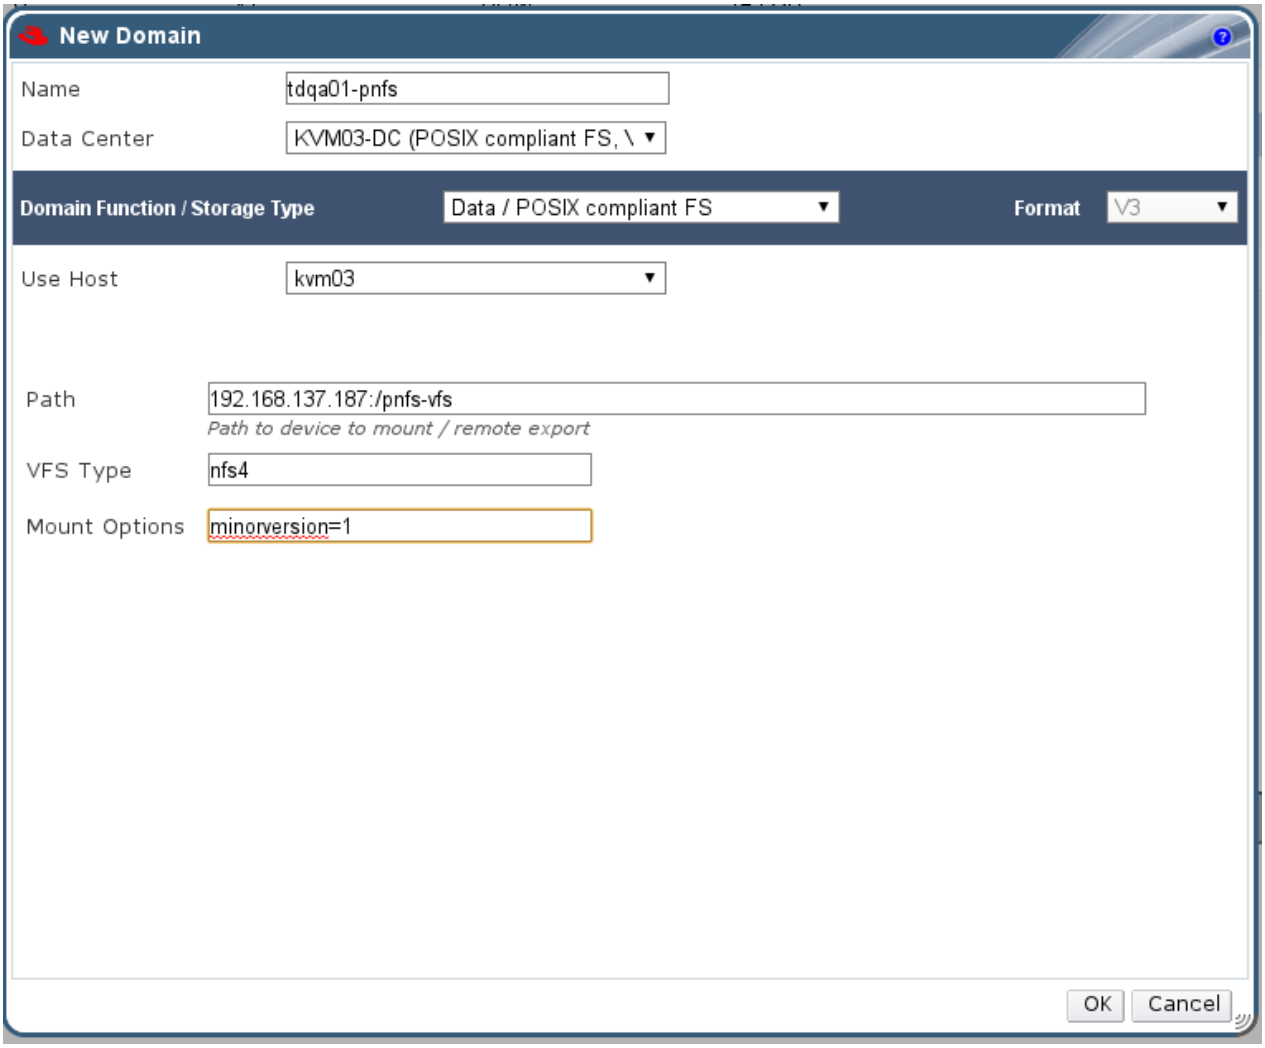

Figure 12.2. NFS Storage

3. Enter the **Name** of the storage domain.
4. Select the **Data Center**, **Domain Function / Storage Type**, and **Use Host** from the drop-down menus.

   If applicable, select the **Format** from the drop-down menu.
5. Enter the **Export Path** to be used for the storage domain.

   The export path should be in the format of **192.168.0.10:/data** or **domain.example.com:/data**
6. In the **VFS Type** field, enter **nfs4**.
7. In the **Mount Options** field, enter **minorversion=1**.

Important

All communication to the storage domain comes from the selected host and not from the Red Hat Enterprise Virtualization Manager. At least one active host must be attached to the chosen Data Center before the storage is configured.

8. Click **OK** to create the storage domain and close the window.

Result

The new pNFS data domain is displayed on the **Storage** tab with a status of **Locked** while the disk prepares. It is automatically attached to the data center upon completion.

Report a bug

12.3.3. Adding iSCSI Storage

Summary

Red Hat Enterprise Virtualization platform supports iSCSI storage by creating a storage domain from a volume group made of pre-existing LUNs. Neither volume groups nor LUNs can be attached to more than one storage domain at a time.

For information regarding the setup and configuration of iSCSI on Red Hat Enterprise Linux, see the *Red Hat Enterprise Linux Storage Administration Guide*.

Note

You can only add an iSCSI storage domain to a data center that is set up for iSCSI storage type.

Procedure 12.4. Adding iSCSI Storage

1. Click the **Storage** resource tab to list the existing storage domains in the results list.
2. Click the **New Domain** button to open the **New Domain** window.
3. Enter the **Name** of the new storage domain.

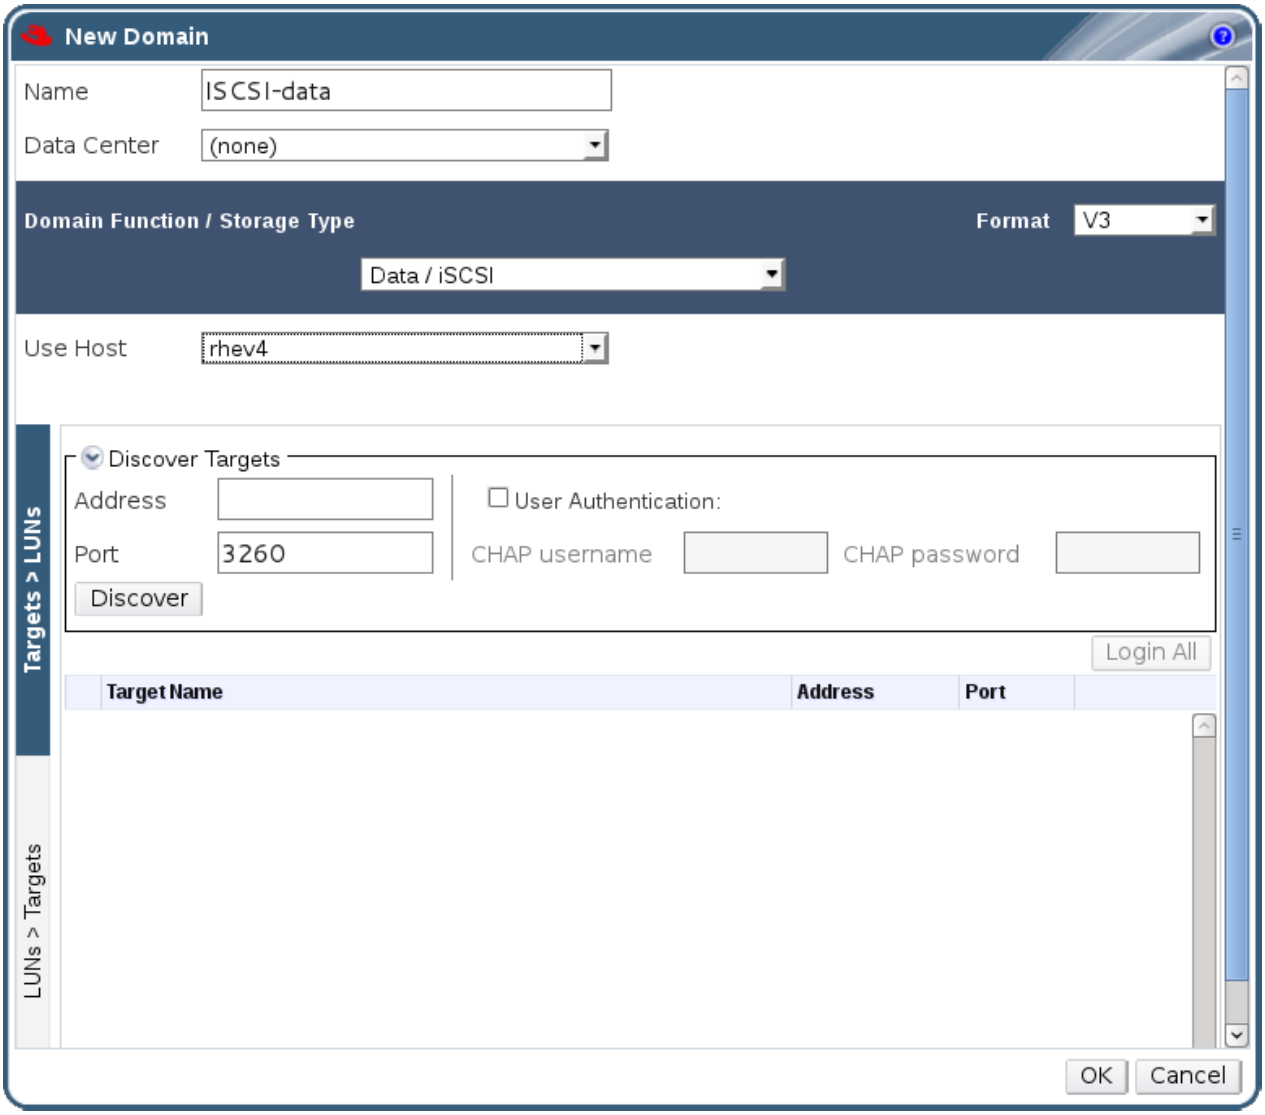

Figure 12.3. New iSCSI Domain

4. Use the **Data Center** drop-down menu to select an iSCSI data center.

   If you do not yet have an appropriate iSCSI data center, select **(none)**.

5. Use the drop-down menus to select the **Domain Function / Storage Type** and the **Format**. The storage domain types that are not compatible with the chosen data center are not available.

6. Select an active host in the **Use Host** field. If this is not the first data domain in a data center, you must select the data center's SPM host.

Important

All communication to the storage domain is via the selected host and not directly from the Red Hat Enterprise Virtualization Manager. At least one active host must exist in the system, and be attached to the chosen data center, before the storage is configured.

7. The Red Hat Enterprise Virtualization Manager is able to map either iSCSI targets to LUNs, or LUNs to iSCSI targets. The **New Domain** window automatically displays known targets with

unused LUNs when iSCSI is selected as the storage type. If the target that you are adding storage from is not listed then you can use target discovery to find it, otherwise proceed to the next step.

iSCSI Target Discovery

a. Click **Discover Targets** to enable target discovery options. When targets have been discovered and logged in to, the **New Domain** window automatically displays targets with LUNs unused by the environment.

Note

LUNs used externally to the environment are also displayed.

You can use the **Discover Targets** options to add LUNs on many targets, or multiple paths to the same LUNs.

b. Enter the fully qualified domain name or IP address of the iSCSI host in the **Address** field.

c. Enter the port to connect to the host on when browsing for targets in the **Port** field. The default is **3260**.

d. If the Challenge Handshake Authentication Protocol (CHAP) is being used to secure the storage, select the **User Authentication** check box. Enter the **CHAP user name** and **CHAP password**.

e. Click the **Discover** button.

f. Select the target to use from the discovery results and click the **Login** button.

Alternatively, click the **Login All** to log in to all of the discovered targets.

8. Click the **+** button next to the desired target. This will expand the entry and display all unused LUNs attached to the target.
9. Select the check box for each LUN that you are using to create the storage domain.
10. Click **OK** to create the storage domain and close the window.

Result

The new iSCSI storage domain displays on the storage tab. This can take up to 5 minutes.

Report a bug

12.3.4. Adding FCP Storage

Summary

Red Hat Enterprise Virtualization platform supports SAN storage by creating a storage domain from a volume group made of pre-existing LUNs. Neither volume groups nor LUNs can be attached to more than one storage domain at a time.

Red Hat Enterprise Virtualization system administrators need a working knowledge of Storage Area Networks (SAN) concepts. SAN usually uses Fibre Channel Protocol (FCP) for traffic between hosts and shared external storage. For this reason, SAN may occasionally be referred to as FCP storage.

For information regarding the setup and configuration of FCP or multipathing on Red Hat Enterprise Linux, please refer to the *Storage Administration Guide* and *DM Multipath Guide*.

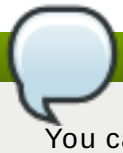

Note

You can only add an FCP storage domain to a data center that is set up for FCP storage type.

Procedure 12.5. Adding FCP Storage

1. Click the **Storage** resource tab to list all storage domains in the virtualized environment.
2. Click **New Domain** to open the **New Domain** window.
3. Enter the **Name** of the storage domain

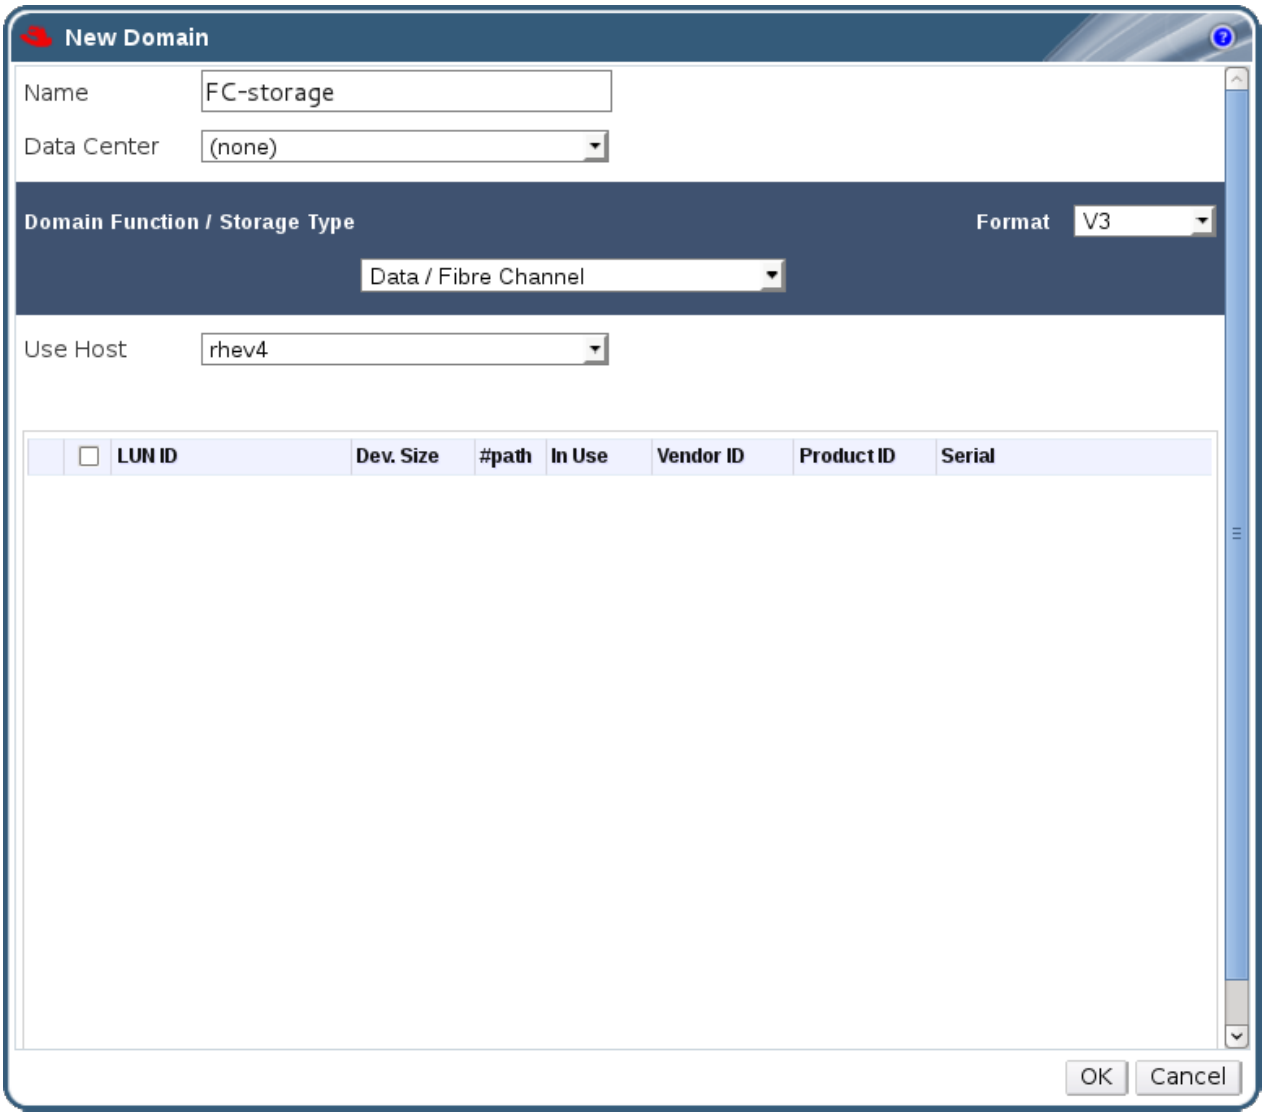

Figure 12.4. Adding FCP Storage

4. Use the **Data Center** drop-down menu to select an FCP data center.

   If you do not yet have an appropriate FCP data center, select **(none)**.

5. Use the drop-down menus to select the **Domain Function / Storage Type** and the **Format**. The storage domain types that are not compatible with the chosen data center are not available.

6. Select an active host in the **Use Host** field. If this is not the first data domain in a data center, you must select the data center's SPM host.

Important

All communication to the storage domain is via the selected host and not directly from the Red Hat Enterprise Virtualization Manager. At least one active host must exist in the system, and be attached to the chosen data center, before the storage is configured.

7. The **New Domain** window automatically displays known targets with unused LUNs when **Data / Fibre Channel** is selected as the storage type. Select the **LUN ID** check box to select all of the available LUNs.

8. Click **OK** to create the storage domain and close the window.

Result

The new FCP data domain displays on the **Storage** tab. It will remain with a **Locked** status while it is being prepared for use. When ready, it is automatically attached to the data center.

Report a bug

12.3.5. Adding Local Storage

12.3.5.1. Preparing Local Storage

Summary

A local storage domain can be set up on a host. When you set up host to use local storage, the host automatically gets added to a new data center and cluster that no other hosts can be added to. Multiple host clusters require that all hosts have access to all storage domains, which is not possible with local storage. Virtual machines created in a single host cluster cannot be migrated, fenced or scheduled.

Important

On Red Hat Enterprise Virtualization Hypervisors the only path permitted for use as local storage is **/data/images**. This directory already exists with the correct permissions on Hypervisor installations. The steps in this procedure are only required when preparing local storage on Red Hat Enterprise Linux virtualization hosts.

Procedure 12.6. Preparing Local Storage

1. On the virtualization host, create the directory to be used for the local storage.

```
# mkdir -p /data/images
```

2. Ensure that the directory has permissions allowing read/write access to the **vdsm** user (UID 36) and **kvm** group (GID 36).

```
# chown 36:36 /data /data/images
```

```
# chmod 0755 /data /data/images
```

Result

Your local storage is ready to be added to the Red Hat Enterprise Virtualization environment.

Report a bug

12.3.5.2. Adding Local Storage

Summary

Storage local to your host has been prepared. Now use the Manager to add it to the host.

Adding local storage to a host in this manner causes the host to be put in a new data center and cluster. The local storage configuration window combines the creation of a data center, a cluster, and storage into a single process.

Procedure 12.7. Adding Local Storage

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click **Maintenance** to place the host into maintenance mode.
3. Click **Configure Local Storage** to open the **Configure Local Storage** window.

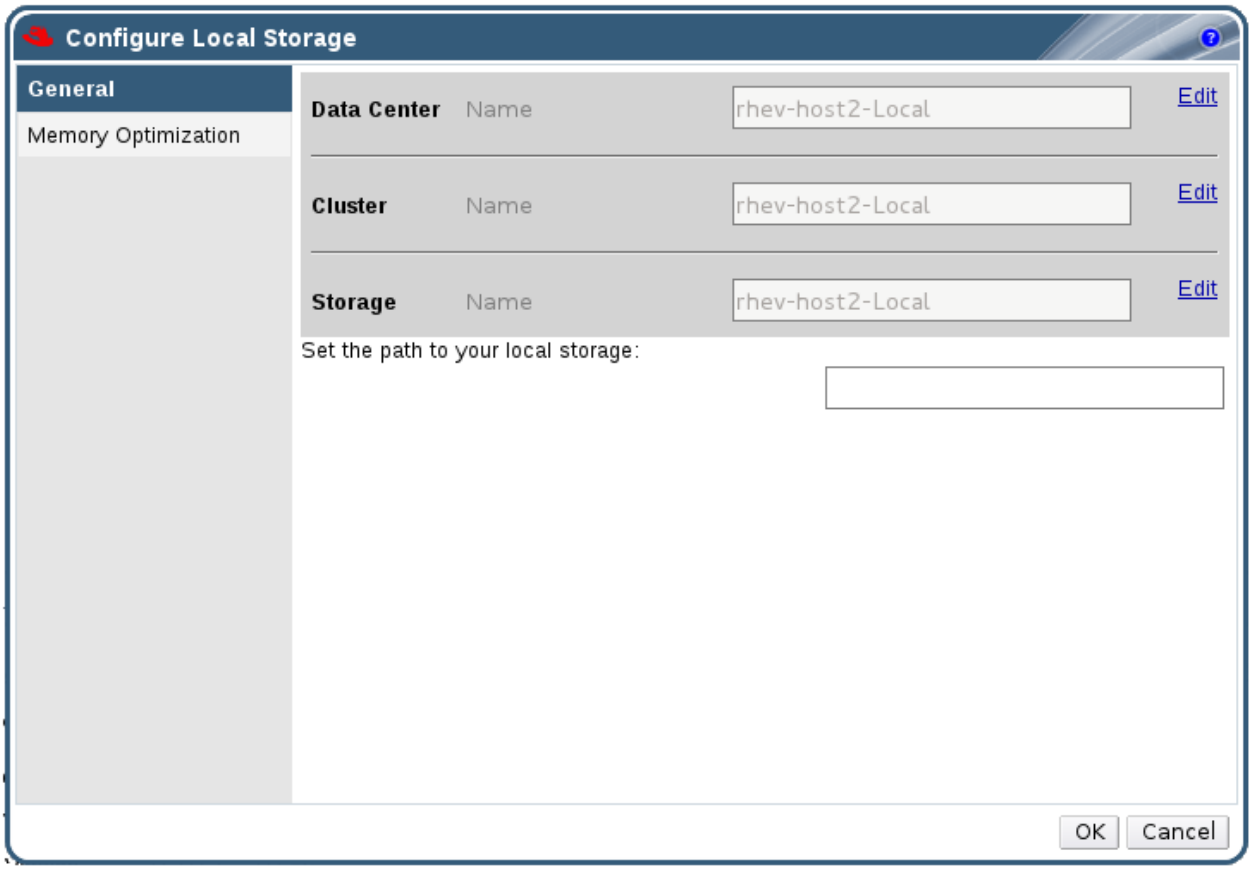

Figure 12.5. Configure Local Storage Window

4. Click the **Edit** buttons next to the **Data Center**, **Cluster**, and **Storage** fields to configure and name the local storage domain.
5. Set the path to your local storage in the text entry field.
6. If applicable, select the **Memory Optimization** tab to configure the memory optimization policy for the new local storage cluster.
7. Click **OK** to save the settings and close the window.

Result

Your host comes online in a data center of its own.

Report a bug

12.3.6. Adding POSIX Compliant File System Storage

12.3.6.1. POSIX Compliant File System Storage in Red Hat Enterprise Virtualization

Red Hat Enterprise Virtualization 3.1 and higher supports the use of POSIX (native) file systems for storage. POSIX file system support allows you to mount file systems using the same mount options that you would normally use when mounting them manually from the command line. This functionality is intended to allow access to storage not exposed using NFS, iSCSI, or FCP.

Any POSIX compliant filesystem used as a storage domain in Red Hat Enterprise Virtualization **MUST** support sparse files and direct I/O. The Common Internet File System (CIFS), for example, does not support direct I/O, making it incompatible with Red Hat Enterprise Virtualization.

Important

Do *not* mount NFS storage by creating a POSIX compliant file system Storage Domain. Always create an NFS Storage Domain instead.

Report a bug

12.3.6.2. Attaching POSIX Compliant File System Storage

Summary

You want to use a POSIX compliant file system that is not exposed using NFS, iSCSI, or FCP as a storage domain.

Procedure 12.8. Attaching POSIX Compliant File System Storage

1. Click the **Storage** resource tab to list the existing storage domains in the results list.
2. Click **New Domain** to open the **New Domain** window.

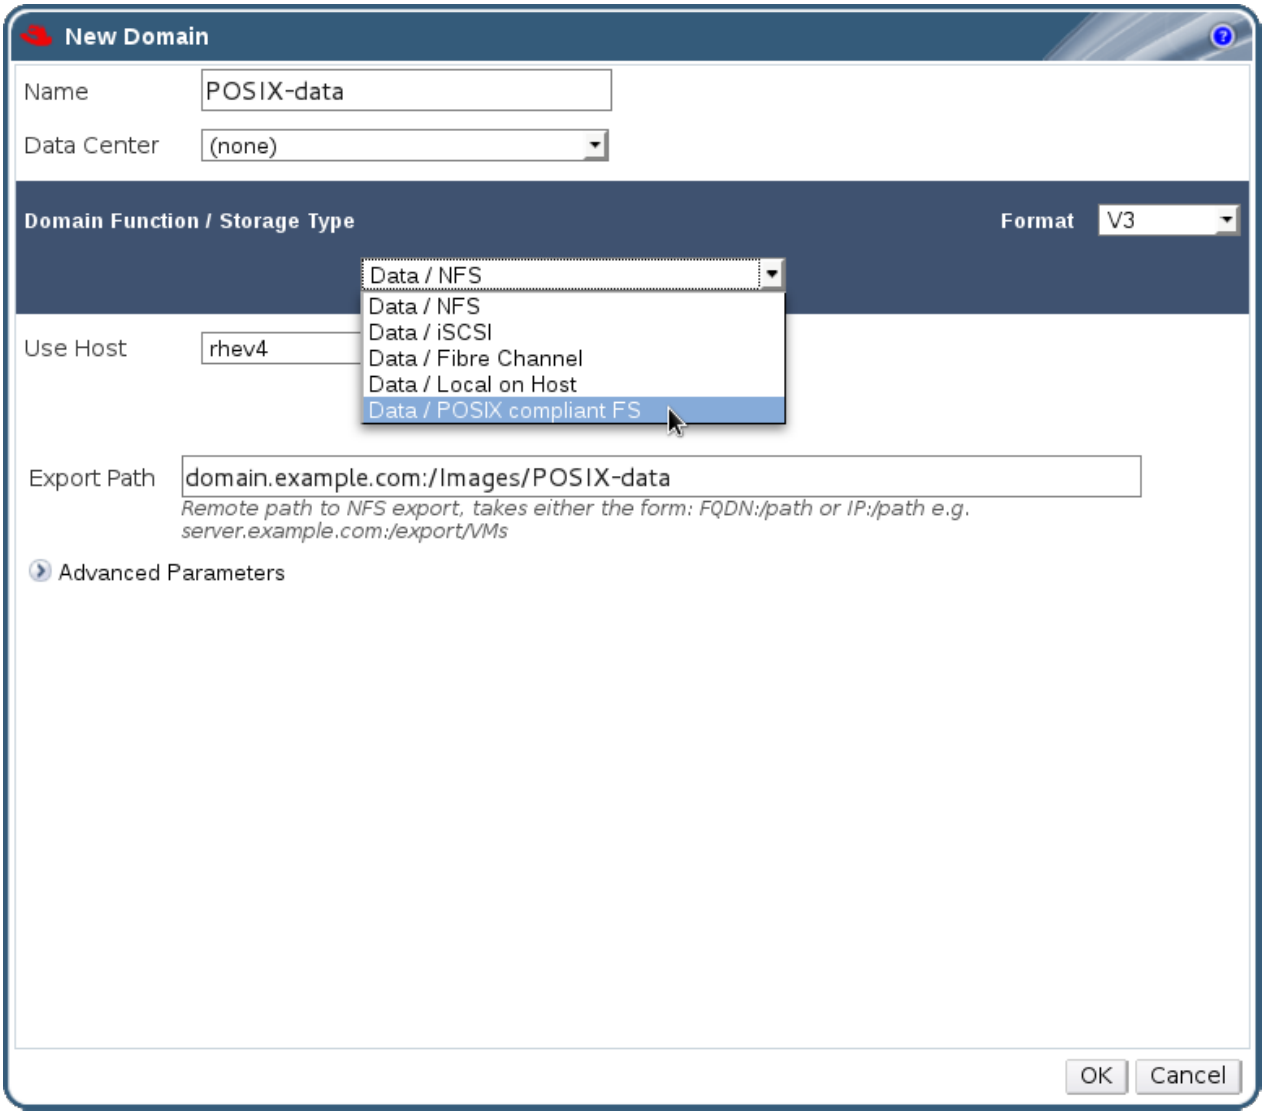

Figure 12.6. POSIX Storage

3. Enter the **Name** for the storage domain.
4. Select the **Data Center** to be associated with the storage domain. The Data Center selected must be of type **POSIX (POSIX compliant FS)**. Alternatively, select **(none)**.
5. Select **Data / POSIX compliant FS** from the **Domain Function / Storage Type** drop-down menu.

   If applicable, select the **Format** from the drop-down menu.
6. Select a host from the **Use Host** drop-down menu. Only hosts within the selected data center will be listed. The host that you select will be used to connect the storage domain.
7. Enter the **Path** to the POSIX file system, as you would normally provide it to the **mount** command.
8. Enter the **VFS Type**, as you would normally provide it to the **mount** command using the ***-t*** argument. See **man mount** for a list of valid VFS types.
9. Enter additional **Mount Options**, as you would normally provide them to the **mount** command using the ***-o*** argument. The mount options should be provided in a comma-separated list. See **man mount** for a list of valid mount options.
10. Click **OK** to attach the new Storage Domain and close the window.

Result

You have used a supported mechanism to attach an unsupported file system as a storage domain.

Report a bug

12.4. Populating the ISO Domain

12.4.1. Populating the ISO Storage Domain

Summary

An ISO storage domain is attached to a data center, ISO images must be uploaded to it. Red Hat Enterprise Virtualization provides an ISO uploader tool that ensures that the images are uploaded into the correct directory path, with the correct user permissions.

The creation of ISO images from physical media is not described in this document. It is assumed that you have access to the images required for your environment.

Procedure 12.9. Populating the ISO Storage Domain

1. Copy the required ISO image to a temporary directory on the system running Red Hat Enterprise Virtualization Manager.
2. Log in to the system running Red Hat Enterprise Virtualization Manager as the **root** user.
3. Use the **engine-iso-uploader** command to upload the ISO image. This action will take some time, the amount of time varies depending on the size of the image being uploaded and available network bandwidth.

   > **Example 12.1. ISO Uploader Usage**

In this example the ISO image **RHEL6.iso** is uploaded to the ISO domain called **ISODomain** using NFS. The command will prompt for an administrative user name and password. The user name must be provided in the form *user name@domain*.

```
# engine-iso-uploader --iso-domain=ISODomain upload RHEL6.iso
```

Result

The ISO image is uploaded and appears in the ISO storage domain specified. It is also available in the list of available boot media when creating virtual machines in the data center which the storage domain is attached to.

See Also:

- Section 12.4.2, “VirtIO and Guest Tool Image Files”
- Section 12.4.3, “Uploading the VirtIO and Guest Tool Image Files to an ISO Storage Domain”

Report a bug

12.4.2. VirtIO and Guest Tool Image Files

The *virtio-win* ISO and Virtual Floppy Drive (VFD) images, which contain the VirtIO drivers for Windows virtual machines, and the *rhev-tools-setup* ISO, which contains the Red Hat Enterprise Virtualization Guest Tools for Windows virtual machines, are copied to an ISO storage domain upon installation and configuration of the domain.

These image files provide software that can be installed on virtual machines to improve performance and usability. The most recent *virtio-win* and *rhev-tools-setup* files can be accessed via the following symbolic links on the file system of the Red Hat Enterprise Virtualization Manager:

- **/usr/share/virtio-win/virtio-win.iso**
- **/usr/share/virtio-win/virtio-win_x86.vfd**
- **/usr/share/virtio-win/virtio-win_amd64.vfd**
- **/usr/share/rhev-guest-tools-iso/rhev-tools-setup.iso**

These image files must be manually uploaded to ISO storage domains that were not created locally by the installation process. Use the **engine-iso-uploader** command to upload these images to your ISO storage domain. Once uploaded, the image files can be attached to and used by virtual machines.

See Also:

- Section 12.4.3, “Uploading the VirtIO and Guest Tool Image Files to an ISO Storage Domain”

Report a bug

12.4.3. Uploading the VirtIO and Guest Tool Image Files to an ISO Storage Domain

The example below demonstrates the command to upload the **virtio-win.iso**, **virtio-win_x86.vfd**, **virtio-win_amd64.vfd**, and **rhev-tools-setup.iso** image files to the **ISODomain**.

Example 12.2. Uploading the VirtIO and Guest Tool Image Files

```
# engine-iso-uploader --iso-domain=[ISODomain] upload
/usr/share/virtio-win/virtio-win.iso /usr/share/virtio-win/virtio-
win_x86.vfd /usr/share/virtio-win/virtio-win_amd64.vfd
/usr/share/rhev-guest-tools-iso/rhev-tools-setup.iso
```

See Also:

- Section 12.4.2, “VirtIO and Guest Tool Image Files”

Report a bug

Log Files

A.1. Red Hat Enterprise Virtualization Manager Installation Log Files

Table A.1. Installation

Log File	Description
/var/log/ovirt-engine/engine-cleanup_*yyyy_mm_dd_hh_mm_ss*.log	Log from the **engine-cleanup** command. This is the command used to reset a Red Hat Enterprise Virtualization Manager installation. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist.
/var/log/ovirt-engine/engine-db-install-*yyyy_mm_dd_hh_mm_ss*.log	Log from the **engine-setup** command detailing the creation and configuration of the **rhevm** database.
/var/log/ovirt-engine/rhevm-dwh-setup-*yyyy_mm_dd_hh_mm_ss*.log	Log from the **rhevm-dwh-setup** command. This is the command used to create the **ovirt_engine_history** database for reporting. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist concurrently.
/var/log/ovirt-engine/ovirt-engine-reports-setup-*yyyy_mm_dd_hh_mm_ss*.log	Log from the **rhevm-reports-setup** command. This is the command used to install the Red Hat Enterprise Virtualization Manager Reports modules. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist concurrently.
/var/log/ovirt-engine/setup/ovirt-engine-setup-*yyyymmddhhmmss*.log	Log from the **engine-setup** command. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist concurrently.

Report a bug

A.2. Red Hat Enterprise Virtualization Manager Log Files

Table A.2. Service Activity

Log File	Description
/var/log/ovirt-engine/engine.log	Reflects all Red Hat Enterprise Virtualization Manager GUI crashes, Active Directory look-ups, Database issues, and other events.
/var/log/ovirt-engine/host-deploy	Log files from hosts deployed from the Red Hat Enterprise Virtualization Manager.
/var/lib/ovirt-engine/setup-history.txt	Tracks the installation and upgrade of packages associated with the Red Hat Enterprise Virtualization Manager.

Report a bug

A.3. Red Hat Enterprise Virtualization Host Log Files

Table A.3.

Log File	Description
/var/log/vdsm/libvirt.log	Log file for **libvirt**.
/var/log/vdsm/spm-lock.log	Log file detailing the host's ability to obtain a lease on the Storage Pool Manager role. The log details when the host has acquired, released, renewed, or failed to renew the lease.
/var/log/vdsm/vdsm.log	Log file for VDSM, the Manager's agent on the virtualization host(s).
/tmp/ovirt-host-deploy-*@DATE@*.log	Host deployment log, copied to engine as /var/log/ovirt-engine/host-deploy/ovirt-*@DATE@*-*@HOST@*-*@CORRELATION_ID@*.log after the host has been successfully deployed.

Report a bug

A.4. Remotely Logging Host Activities

A.4.1. Setting Up a Virtualization Host Logging Server

Summary

Red Hat Enterprise Virtualization hosts generate and update log files, recording their actions and problems. Collecting these log files centrally simplifies debugging.

This procedure should be used on your centralized log server. You could use a separate logging server, or use this procedure to enable host logging on the Red Hat Enterprise Virtualization Manager.

Procedure A.1. Setting up a Virtualization Host Logging Server

1. Configure SELinux to allow **rsyslog** traffic.

```
# semanage port -a -t syslogd_port_t -p udp 514
```

2. Edit **/etc/rsyslog.conf** and add below lines:

```
$template TmplAuth, "/var/log/%fromhost%/secure"
$template TmplMsg, "/var/log/%fromhost%/messages"

$RuleSet remote
authpriv.*   ?TmplAuth
*.info,mail.none;authpriv.none,cron.none   ?TmplMsg
$RuleSet RSYSLOG_DefaultRuleset
$InputUDPServerBindRuleset remote
```

Uncomment the following:

```
#$ModLoad imudp
#$UDPServerRun 514
```

3. Restart the rsyslog service:

```
# service rsyslog restart
```

Result

Your centralized log server is now configured to receive and store the **messages** and **secure** logs from your virtualization hosts.

Report a bug

A.4.2. Configuring Logging

Prerequisites:

- Section A.4.1, “Setting Up a Virtualization Host Logging Server”

Summary

The Hypervisor creates and updates a number of log files. The **Logging** screen allows configuration of a daemon to automatically export these log files to a remote server.

Procedure A.2. Configuring Hypervisor Logging

1. **logrotate Configuration**

   The **logrotate** utility simplifies the administration of log files. The Hypervisor uses **logrotate** to rotate logs when they reach a certain file size.

   Log rotation involves renaming the current logs and starting new ones in their place. The **Logrotate Max Log Size** value set on the **Logging** screen is used to determine when a log will be rotated.

   Enter the **Logrotate Max Log Size** in kilobytes. The default maximum log size is 1024 kilobytes.

2. **rsyslog Configuration**

   The **rsyslog** utility is a multithreaded syslog daemon. The Hypervisor is able to use **rsyslog** to transmit log files over the network to a remote syslog daemon. For information on setting up the remote syslog daemon, see the *Red Hat Enterprise Linux Deployment Guide*.

   a. Enter the remote **rsyslog** server address in the **Server Address** field.

   b. Enter the remote **rsyslog** server port in the **Server Port** field. The default port is **514**.

3. **netconsole Configuration**

   The **netconsole** module allows kernel messages to be sent to a remote machine. The Hypervisor uses **netconsole** to transmit kernel messages over the network.

   a. Enter the **Server Address**.

b. Enter the **Server Port**. The default port is **6666**.

4. **Save Configuration**

   To save the logging configuration, select **<Save>** and press **Enter**.

Result

The logging configuration has been updated and logs will be exported to the remote **rsyslog** server specified.

Report a bug

A.4.3. Configuring Logging

Prerequisites:

- Section A.4.1, “Setting Up a Virtualization Host Logging Server”

Summary

The Hypervisor creates and updates a number of log files. The **Logging** screen allows configuration of a daemon to automatically export these log files to a remote server.

Procedure A.3. Configuring Hypervisor Logging

1. **logrotate Configuration**

   The **logrotate** utility simplifies the administration of log files. The Hypervisor uses **logrotate** to rotate logs when they reach a certain file size.

   Log rotation involves renaming the current logs and starting new ones in their place. The **Logrotate Max Log Size** value set on the **Logging** screen is used to determine when a log will be rotated.

   Enter the **Logrotate Max Log Size** in kilobytes. The default maximum log size is 1024 kilobytes.

2. **rsyslog Configuration**

   The **rsyslog** utility is a multithreaded syslog daemon. The Hypervisor is able to use **rsyslog** to transmit log files over the network to a remote syslog daemon. For information on setting up the remote syslog daemon, see the *Red Hat Enterprise Linux Deployment Guide*.

   a. Enter the remote **rsyslog** server address in the **Server Address** field.

   b. Enter the remote **rsyslog** server port in the **Server Port** field. The default port is **514**.

3. **netconsole Configuration**

   The **netconsole** module allows kernel messages to be sent to a remote machine. The Hypervisor uses **netconsole** to transmit kernel messages over the network.

   a. Enter the **Server Address**.

   b. Enter the **Server Port**. The default port is **6666**.

4. **Save Configuration**

   To save the logging configuration, select **<Save>** and press **Enter**.

Result

The logging configuration has been updated and logs will be exported to the remote **rsyslog** server specified.

Report a bug

Additional Utilities

B.1. Domain Management Tool

B.1.1. What is the Domain Management Tool?

Red Hat Enterprise Virtualization Manager authenticates users using directory services. To add users to Red Hat Enterprise Virtualization Manager you must first use the internal **admin** user to add the directory service that the users must be authenticated against. You add and remove directory services domains using the included domain management tool, **engine-manage-domains**.

The **engine-manage-domains** command is only accessible on the machine on which Red Hat Enterprise Virtualization Manager is installed. The **engine-manage-domains** command must be run as the **root** user.

See Also:

- Section B.1.2, “Syntax for the Domain Management Tool”
- Section B.1.3, “Adding Domains to Configuration”
- Section B.1.4, “Editing a Domain in the Configuration”
- Section B.1.5, “Deleting a Domain from the Configuration”
- Section B.1.6, “Validating Domain Configuration”
- Section B.1.7, “Listing Domains in Configuration”

Report a bug

B.1.2. Syntax for the Domain Management Tool

The usage syntax is:

```
engine-manage-domains -action=ACTION [options]
```

Available actions are:

add

Add a domain to Red Hat Enterprise Virtualization Manager's directory services configuration.

edit

Edit a domain in Red Hat Enterprise Virtualization Manager's directory services configuration.

delete

Delete a domain from Red Hat Enterprise Virtualization Manager's directory services configuration.

validate

Validate Red Hat Enterprise Virtualization Manager's directory services configuration. This

command attempts to authenticate each domain in the configuration using the configured user name and password.

list

List Red Hat Enterprise Virtualization Manager's current directory services configuration.

These options can be combined with the actions on the command line:

-domain=DOMAIN

Specifies the domain on which the action will be performed. The ***-domain*** parameter is mandatory for ***add***, ***edit***, and ***delete***.

***-provider*=PROVIDER**

Specifies the LDAP provider type of the directory server for the domain. Valid values are:

- ***ActiveDirectory*** - Active Directory.
- ***IPA*** - Identity Management (IdM).
- ***RHDS*** - Red Hat Directory Server. Red Hat Directory Server does not come with Kerberos. Red Hat Enterprise Virtualization requires Kerberos authentication. RHDS must be made a service within a Kerberos domain to provide directory services to the Manager.

Note

If you want to use RHDS as your directory server, you must have the **memberof** plugin installed in RHDS. To use the **memberof** plugin, your users must be **inetuser**s. For more information about using the **memberof** plugin, see the *Red Hat Directory Server Plug-in Guide*.

-user=USER

Specifies the domain user to use. The ***-user*** parameter is mandatory for ***add***, and optional for ***edit***.

-passwordFile=FILE

Specifies that the domain user's password is on the first line of the provided file. This option, or the ***-interactive*** option, must be used to provide the password for use with the ***add*** action.

-addPermissions

Specifies that the domain user will be given the **SuperUser** role in Red Hat Enterprise Virtualization Manager. By default, if the ***-addPermissions*** parameter is not specified, the **SuperUser** role is not assigned to the domain user. The ***-addPermissions*** option is optional. It is only valid when used in combination with the ***add*** and ***edit*** actions.

-interactive

Specifies that the domain user's password is to be provided interactively. This option, or the ***-passwordFile*** option, must be used to provide the password for use with the ***add*** action.

-configFile=FILE

Specifies an alternate configuration file that the command must load. The ***-configFile*** parameter is always optional.

-report

In conjunction with the ***validate*** action results in the output of a report of all encountered validation errors.

For full usage information consult the **engine-manage-domains** command's help output:

```
# engine-manage-domains --help
```

See Also:

- Section B.1.1, “What is the Domain Management Tool?”
- Section B.1.3, “Adding Domains to Configuration”
- Section B.1.4, “Editing a Domain in the Configuration”
- Section B.1.5, “Deleting a Domain from the Configuration”
- Section B.1.6, “Validating Domain Configuration”
- Section B.1.7, “Listing Domains in Configuration”

Report a bug

B.1.3. Adding Domains to Configuration

In this example, the **engine-manage-domains** command is used to add the **IdM** domain **directory.demo.redhat.com** to the Red Hat Enterprise Virtualization Manager configuration. The configuration is set to use the **admin** user when querying the domain; the password is provided interactively.

Example B.1. engine-manage-domains Add Action

```
# engine-manage-domains -action=add -domain=directory.demo.redhat.com
-provider=IPA -user=admin -interactive
loaded template kr5.conf file
setting default_tkt_enctypes
setting realms
setting domain realm
success
User guid is: 80b71bae-98a1-11e0-8f20-525400866c73
Successfully added domain directory.demo.redhat.com. oVirt Engine
restart is required in order for the changes to take place (service
ovirt-engine restart).
```

See Also:

- Section B.1.1, “What is the Domain Management Tool?”
- Section B.1.2, “Syntax for the Domain Management Tool”

- Section B.1.4, “Editing a Domain in the Configuration”
- Section B.1.5, “Deleting a Domain from the Configuration”
- Section B.1.6, “Validating Domain Configuration”
- Section B.1.7, “Listing Domains in Configuration”

Report a bug

B.1.4. Editing a Domain in the Configuration

In this example, the **engine-manage-domains** command is used to edit the **directory.demo.redhat.com** domain in the Red Hat Enterprise Virtualization Manager configuration. The configuration is updated to use the **admin** user when querying this domain; the password is provided interactively.

Example B.2. engine-manage-domains Edit Action

```
# engine-manage-domains -action=edit -domain=directory.demo.redhat.com
-user=admin -interactive
loaded template kr5.conf file
setting default_tkt_enctypes
setting realms
setting domain realmo
success
User guide is: 80b71bae-98a1-11e0-8f20-525400866c73
Successfully edited domain directory.demo.redhat.com. oVirt Engine
restart is required in order for the changes to take place (service
ovirt-engine restart).
```

See Also:

- Section B.1.1, “What is the Domain Management Tool?”
- Section B.1.2, “Syntax for the Domain Management Tool”
- Section B.1.3, “Adding Domains to Configuration”
- Section B.1.5, “Deleting a Domain from the Configuration”
- Section B.1.6, “Validating Domain Configuration”
- Section B.1.7, “Listing Domains in Configuration”

Report a bug

B.1.5. Deleting a Domain from the Configuration

In this example, the **engine-manage-domains** command is used to remove the **directory.demo.redhat.com** domain from the Red Hat Enterprise Virtualization Manager configuration. Users defined in the removed domain will no longer be able to authenticate with the Red Hat Enterprise Virtualization Manager. The entries for the affected users will remain defined in the Red Hat Enterprise Virtualization Manager until they are explicitly removed.

The domain being removed in this example is the last one listed in the Red Hat Enterprise Virtualization Manager configuration. A warning is displayed highlighting this fact and that only the **admin** user from the **internal** domain will be able to log in until another domain is added.

Example B.3. engine-manage-domains Delete Action

```
# engine-manage-domains -action=delete -
domain='directory.demo.redhat.com'
WARNING: Domain directory.demo.redhat.com is the last domain in the
configuration. After deleting it you will have to either add another
domain, or to use the internal admin user in order to login.
Successfully deleted domain directory.demo.redhat.com. Please remove
all users and groups of this domain using the Administration portal or
the API.
```

See Also:

- Section B.1.1, “What is the Domain Management Tool?”
- Section B.1.2, “Syntax for the Domain Management Tool”
- Section B.1.3, “Adding Domains to Configuration”
- Section B.1.4, “Editing a Domain in the Configuration”
- Section B.1.6, “Validating Domain Configuration”
- Section B.1.7, “Listing Domains in Configuration”

Report a bug

B.1.6. Validating Domain Configuration

In this example, the **engine-manage-domains** command is used to validate the Red Hat Enterprise Virtualization Manager configuration. The command attempts to log into each listed domain with the credentials provided in the configuration. The domain is reported as valid if the attempt is successful.

Example B.4. engine-manage-domains Validate Action

```
# engine-manage-domains -action=validate
User guide is: 80b71bae-98a1-11e0-8f20-525400866c73
Domain directory.demo.redhat.com is valid.
```

See Also:

- Section B.1.1, “What is the Domain Management Tool?”
- Section B.1.2, “Syntax for the Domain Management Tool”
- Section B.1.3, “Adding Domains to Configuration”
- Section B.1.4, “Editing a Domain in the Configuration”

- Section B.1.5, “Deleting a Domain from the Configuration”
- Section B.1.7, “Listing Domains in Configuration”

Report a bug

B.1.7. Listing Domains in Configuration

The **engine-manage-domains** command lists the directory services domains defined in the Red Hat Enterprise Virtualization Manager configuration. This command prints the domain, the user name in User Principal Name (UPN) format, and whether the domain is local or remote for each configuration entry.

Example B.5. engine-manage-domains List Action

```
# engine-manage-domains -action=list
Domain: directory.demo.redhat.com
    User name: admin@DIRECTORY.DEMO.REDHAT.COM
    This domain is a remote domain.
```

See Also:

- Section B.1.1, “What is the Domain Management Tool?”
- Section B.1.2, “Syntax for the Domain Management Tool”
- Section B.1.3, “Adding Domains to Configuration”
- Section B.1.4, “Editing a Domain in the Configuration”
- Section B.1.5, “Deleting a Domain from the Configuration”
- Section B.1.6, “Validating Domain Configuration”

Report a bug

B.2. Configuration Tool

B.2.1. Configuration Tool

Installing the Red Hat Enterprise Virtualization Manager modifies only a subset of configuration settings from their defaults. Further modifications are made using the configuration tool: **engine-config**.

The configuration tool does not require Red Hat JBoss Enterprise Application Platform or the Red Hat Enterprise Virtualization Manager to be running to update the configuration. Configuration key values are stored in the database; configuration changes will not be saved unless the database is operational. Changes are applied when Red Hat JBoss Enterprise Application Platform is restarted.

The Red Hat Enterprise Virtualization Manager stores configuration settings as a series of key-to-value pair mappings. The configuration tool allows you to:

- list all available configuration keys,
- list all available configuration values,

- retrieve the value of a specific configuration key, and
- set the value of a specific configuration key.

You are also able to maintain multiple versions of the Manager's configuration with the configuration tool. Use the ***--cver*** parameter to specify the configuration version to be used when retrieving or setting a value for a configuration key. The default configuration version is **general**.

See Also:

- Section B.2.2, "Syntax for engine-config Command"

Report a bug

B.2.2. Syntax for engine-config Command

The configuration tool is accessible on the client machine on which the Red Hat Enterprise Virtualization Manager is installed. For full usage information consult the help output of the **engine-config** command:

```
# engine-config --help
```

Common tasks

List available configuration keys

Use the ***--list*** parameter to list available configuration keys.

```
# engine-config --list
```

Each available configuration key is listed by name and description.

List available configuration values

Use the ***--all*** parameter to list available configuration values.

```
# engine-config --all
```

Each available configuration key is listed by name, current value of the key, and the configuration version.

Retrieve value of configuration key

Use the ***--get*** parameter to retrieve the value of a specific key.

```
# engine-config --get KEY_NAME
```

Replace *KEY_NAME* with the name of the specific key to retrieve the key name, value, and the configuration version. Use the ***--cver*** parameter to specify the configuration version of the value to be retrieved.

Set value of configuration key

Use the ***--set*** parameter to set the value of a specific key. You must also set the configuration version to which the change is to apply using the ***--cver*** parameter.

```
# engine-config --set KEY_NAME=KEY_VALUE --cver=VERSION
```

Replace *KEY_NAME* with the name of the specific key to set; replace *KEY_VALUE* with the value to be set. Environments with more than one configuration version require the *VERSION* to be specified.

See Also:

- Section B.2.1, “Configuration Tool”

Report a bug

B.3. Image Uploader

B.3.1. Virtual Machine Image Uploader

Using the **engine-image-uploader** command, you can list export storage domains and upload virtual machines in OVF to an export storage domain and have them automatically recognized in the Red Hat Enterprise Virtualization Manager. The tool only supports gzip compressed OVF files created by Red Hat Enterprise Virtualization.

The image uploader makes creating distributable virtual machine images practical.

The archive contains images and master directories that are in the following format:

```
|-- images
|   |-- [Image Group UUID]
|         |--- [Image UUID (this is the disk image)]
|         |--- [Image UUID (this is the disk image)].meta
|-- master
|   |---vms
|       |--- [UUID]
|             |--- [UUID].ovf
```

See Also:

- Section B.3.2, “Syntax for the engine-image-uploader Command”
- Section B.3.3, “Creating an OVF Archive That is Compatible with the Image Uploader”
- Section B.3.4, “Basic engine-image-uploader Usage Examples”

Report a bug

B.3.2. Syntax for the engine-image-uploader Command

The basic syntax for the image uploader command is:

```
engine-image-uploader [options] list
engine-image-uploader [options] upload [file].[file]...[file]
```

The two supported modes of operation are ***list*** and ***upload***.

- The ***list*** parameter lists the valid export storage domains available for image uploads.

- The ***upload*** parameter uploads selected image file(s) to the specified image storage domain.

The image uploader command requires either the ***list*** or ***upload*** parameter be included for basic usage. The ***upload*** parameter requires a minimum of one local file name to upload.

There are numerous parameters to further refine the **engine-image-uploader** command. You can set defaults for any of these in the **/etc/ovirt-engine/imageuploader.conf** file.

General Options

-h*, *--help

Displays command usage information and returns to prompt.

--conf-file=PATH

Sets *PATH* as the configuration file the tool is to use. The default is **etc/ovirt-engine/imageuploader.conf**.

--log-file=PATH

Sets *PATH* as the specific file name the command should use for the log output.

--quiet

Sets quiet mode, reducing console output to a minimum. Quiet mode is off by default.

-v*, *--verbose

Sets verbose mode, providing more console output. Verbose mode is off by default.

-f*, *--force

Force mode is necessary when the source file being uploaded has an identical file name as an existing file at the destination; it forces the existing file to be overwritten. Force mode is off by default.

Red Hat Enterprise Virtualization Manager Options

-u USER*, *--user=USER

Sets the user associated with the file to be uploaded. The *USER* is specified in the format *user@domain*, where *user* is the user name and *domain* is the directory services domain in use. The user must exist in directory services and be known to the Red Hat Enterprise Virtualization Manager.

-r FQDN*, *--rhevm=FQDN

Sets the fully qualified domain name of the Red Hat Enterprise Virtualization Manager server from which to upload images, where *FQDN* is replaced by the fully qualified domain name of the Manager. It is assumed that the image uploader is being run on the same client machine as the Red Hat Enterprise Virtualization Manager; the default value is **localhost:443**.

Export Storage Domain Options

These options specify the export domain to which files are uploaded. They are alternatives; do not use these parameters together.

-e*, *--export-domain=EXPORT_DOMAIN

Sets the storage domain *EXPORT_DOMAIN* as the destination for uploads.

-n*, *--nfs-server=NFSSERVER

Sets the NFS path *NFSSERVER* as the destination for uploads.

-i*, *--ovf-id

Use this option if you do not want to update the UUID of the image. By default, the tool will generate a new UUID for the image. This ensures that there is no conflict between the id of the incoming image and those already in the environment.

-d*, *-disk-instance-id

Use this option if you do not want to rename the instance ID for each disk (i.e. InstanceId) in the image. By default, this tool will generate new UUIDs for disks within the image to be imported. This ensures that there are no conflicts between the disks on the imported image and those within the environment.

-m*, *--mac-address

Use this option if you do not want to remove the network components from the image that will be imported. By default, this tool will remove any network interface cards from the image to prevent conflicts with network cards on other virtual machines within the environment. Once the image has been imported, use the Administration Portal to add network interface cards back and the Manager will ensure that there are no MAC address conflicts.

-N NEW_IMAGE_NAME*, *--name=NEW_IMAGE_NAME

Supply this option if you want to rename the image.

See Also:

- Section B.3.1, “Virtual Machine Image Uploader”
- Section B.3.3, “Creating an OVF Archive That is Compatible with the Image Uploader”
- Section B.3.4, “Basic engine-image-uploader Usage Examples”

Report a bug

B.3.3. Creating an OVF Archive That is Compatible with the Image Uploader

Summary

You can create files that can be uploaded using the **`engine-image-uploader`** tool.

Procedure B.1. Creating an OVF Archive That is Compatible with the Image Uploader

1. Use the Manager to create an empty export domain. An empty export domain makes it easy to see which directory contains your virtual machine.
2. Export your virtual machine to the empty export domain you just created.
3. Log in to the storage server that serves as the export domain, find the root of the NFS share and change to the subdirectory under that mount point. You started with a new export domain, there is only one directory under the exported directory. It contains the **`images/`** and **`master/`** directories.

4. Run the **tar -zcvf my.ovf images/ master/** command to create the tar/gzip ovf archive.

5. Anyone you give the resulting ovf file to (in this example, called **my.ovf**) can import it to Red Hat Enterprise Virtualization Manager using the **engine-image-uploader** command.

Result

You have created a compressed OVF image file that can be distributed. Anyone you give it to can use the **engine-image-uploader** command to upload your image into their Red Hat Enterprise Virtualization environment.

See Also:

- Section B.3.1, “Virtual Machine Image Uploader”
- Section B.3.2, “Syntax for the engine-image-uploader Command”
- Section B.3.4, “Basic engine-image-uploader Usage Examples”

Report a bug

B.3.4. Basic engine-image-uploader Usage Examples

Here is an example of how to use **engine-image-uploader** to list storage domains:

Example B.6. Uploading a file Using the engine-image-uploader Tool

```
# engine-image-uploader list
Please provide the REST API username for RHEV-M: admin@internal
Please provide the REST API password for the admin@internal RHEV-M
user: **********
Export Storage Domain Name | Datacenter  | Export Domain Status
myexportdom                | Myowndc     | active
```

To upload an Open Virtualization Format (ovf) file, you need to enter an NFS server name (**-n *NFSSERVER***) or export domain (**-e *EXPORT_STORAGE_DOMAIN***) and the name of the **.ovf** file:

```
# engine-image-uploader -e myexportdom upload myrhel6.ovf
Please provide the REST API username for RHEV-M: admin@internal
Please provide the REST API password for the admin@internal RHEV-M
user: **********
```

See Also:

- Section B.3.1, “Virtual Machine Image Uploader”
- Section B.3.2, “Syntax for the engine-image-uploader Command”
- Section B.3.3, “Creating an OVF Archive That is Compatible with the Image Uploader”

Report a bug

B.4. ISO Uploader

B.4.1. ISO Uploader

The ISO uploader is a tool for uploading ISO images to the ISO storage domain. It is installed as part of the Red Hat Enterprise Virtualization Manager.

The ISO uploader command is **engine-iso-uploader**. You are required to log in as the **root** user and provide the administration credentials for the Red Hat Enterprise Virtualization environment. The **engine-iso-uploader *-h*** command displays usage information, including a list of all valid options for the **engine-iso-uploader** command.

See Also:

- Section B.4.2, “Syntax for engine-iso-uploader Command”
- Section B.4.3.1, “Specifying an NFS Server”
- Section B.4.3.2, “Basic ISO Uploader Usage”

Report a bug

B.4.2. Syntax for engine-iso-uploader Command

The basic syntax for the ISO uploader command is:

```
engine-iso-uploader [options] list
engine-iso-uploader [options] upload [file].[file]...[file]
```

The two supported modes of operation are ***list*** and ***upload***.

- The ***list*** parameter lists the valid ISO storage domains available for ISO uploads. The Red Hat Enterprise Virtualization Manager sets this list on the local machine upon installation.
- The ***upload*** parameter uploads single or multiple space-separated ISO files to the specified ISO storage domain. NFS is used as default; SSH is available.

The ISO uploader command requires either the ***list*** or ***upload*** parameter be included for basic usage. The ***upload*** parameter requires a minimum of one local file name to upload.

There are numerous parameters to further refine the **engine-iso-uploader** command.

General Options

--version

Displays the version number of the command in use and returns to prompt.

-h*, *--help

Displays command usage information and returns to prompt.

--conf-file=PATH

Sets *PATH* as the configuration file the tool is to use.

--log-file=PATH

Sets *PATH* as the specific file name the command should use for the log output.

--quiet

Sets quiet mode, reducing console output to a minimum. Quiet mode is off by default.

-v*, *--verbose

Sets verbose mode, providing more console output. Verbose mode is off by default.

-f*, *--force

Force mode is necessary when the source file being uploaded has an identical file name as an existing file at the destination; it forces the existing file to be overwritten. Force mode is off by default.

Red Hat Enterprise Virtualization Manager Options

-u USER*, *--user=USER

Sets the user associated with the file to be uploaded. The *USER* is specified in the format *user@domain*, where *user* is the user name and *domain* is the directory services domain in use. The user must exist in directory services and be known to the Red Hat Enterprise Virtualization Manager.

-r FQDN*, *--rhevm=FQDN

Sets the fully qualified domain name of the Red Hat Enterprise Virtualization Manager server from which to upload ISOs, where *FQDN* is replaced by the fully qualified domain name of the Manager. It is assumed that the ISO uploader is being run on the same client machine as the Red Hat Enterprise Virtualization Manager; the default value is **`localhost`**.

ISO Storage Domain Options

These options specify the ISO domain to which files are uploaded. They are alternatives; do not use these parameters together.

-i*, *--iso-domain=ISODOMAIN

Sets the storage domain *ISODOMAIN* as the destination for uploads.

-n*, *--nfs-server=NFSSERVER

Sets the NFS path *NFSSERVER* as the destination for uploads.

Connection Options

The ISO uploader uses NFS as default to upload files. These options specify SSH file transfer instead.

--ssh-user=USER

Sets *USER* as the SSH user name to use for the upload.

--ssh-port=PORT

Sets *PORT* as the port to use when connecting to SSH.

-k KEYFILE*, *--key-file=KEYFILE

Sets *KEYFILE* as the public key to use for SSH authentication. You will be prompted to enter the password of the user specified with ***--ssh-user=USER*** if no key is set.

See Also:

- Section B.4.1, “ISO Uploader”
- Section B.4.3.1, “Specifying an NFS Server”
- Section B.4.3.2, “Basic ISO Uploader Usage”

Report a bug

B.4.3. Usage Examples

B.4.3.1. Specifying an NFS Server

Example B.7. Uploading to an NFS Server

```
# engine-iso-uploader --nfs-server=storage.demo.redhat.com:/iso/path
upload RHEL6.0.iso
```

See Also:

- Section B.4.1, “ISO Uploader”
- Section B.4.2, “Syntax for engine-iso-uploader Command”
- Section B.4.3.2, “Basic ISO Uploader Usage”

Report a bug

B.4.3.2. Basic ISO Uploader Usage

The example below demonstrates the ISO uploader and the list parameter. The first command lists the available ISO storage domains; the **admin@internal** user is used because no user was specified in the command. The second command uploads an ISO file over NFS to the specified ISO domain.

Example B.8. List Domains and Upload Image

```
# engine-iso-uploader list
Please provide the REST API password for the admin@internal oVirt
Engine user (CTRL+D to abort):
ISO Storage Domain Name   | Datacenter          | ISO Domain Status
ISODomain                 | Default             | active
```

```
# engine-iso-uploader --iso-domain=[ISODomain] upload [RHEL6.iso]
Please provide the REST API password for the admin@internal oVirt
Engine user (CTRL+D to abort):
```

See Also:

- Section B.4.1, “ISO Uploader”
- Section B.4.2, “Syntax for engine-iso-uploader Command”

- Section B.4.3.1, “Specifying an NFS Server”

Report a bug

B.5. Log Collector

B.5.1. Log Collector

A log collection tool is included in the Red Hat Enterprise Virtualization Manager. This allows you to easily collect relevant logs from across the Red Hat Enterprise Virtualization environment when requesting support.

The log collection command is **engine-log-collector**. You are required to log in as the **root** user and provide the administration credentials for the Red Hat Enterprise Virtualization environment. The **engine-log-collector *-h*** command displays usage information, including a list of all valid options for the **engine-log-collector** command.

See Also:

- Section B.5.2, “Syntax for engine-log-collector Command”
- Section B.5.3, “Basic Log Collector Usage”
- Appendix A, *Log Files*

Report a bug

B.5.2. Syntax for engine-log-collector Command

The basic syntax for the log collector command is:

```
engine-log-collector [options] list [all, clusters, datacenters]
engine-log-collector [options] collect
```

The two supported modes of operation are ***list*** and ***collect***.

- The ***list*** parameter lists either the hosts, clusters, or data centers attached to the Red Hat Enterprise Virtualization Manager. You are able to filter the log collection based on the listed objects.
- The ***collect*** parameter performs log collection from the Red Hat Enterprise Virtualization Manager. The collected logs are placed in an archive file under the **/tmp/logcollector** directory. The **engine-log-collector** command assigns each log a specific file name.

Unless another parameter is specified, the default action is to list the available hosts together with the data center and cluster to which they belong. You will be prompted to enter user names and passwords to retrieve certain logs.

There are numerous parameters to further refine the **engine-log-collector** command.

General options

--version

Displays the version number of the command in use and returns to prompt.

-h*, *--help

Displays command usage information and returns to prompt.

--conf-file=PATH

Sets *PATH* as the configuration file the tool is to use.

--local-tmp=PATH

Sets *PATH* as the directory in which logs are saved. The default directory is **/tmp/logcollector**.

--ticket-number=TICKET

Sets *TICKET* as the ticket, or case number, to associate with the SOS report.

--upload=FTP_SERVER

Sets *FTP_SERVER* as the destination for retrieved logs to be sent using FTP. Do not use this option unless advised to by a Red Hat support representative.

--log-file=PATH

Sets *PATH* as the specific file name the command should use for the log output.

--quiet

Sets quiet mode, reducing console output to a minimum. Quiet mode is off by default.

-v*, *--verbose

Sets verbose mode, providing more console output. Verbose mode is off by default.

Red Hat Enterprise Virtualization Manager Options

These options filter the log collection and specify authentication details for the Red Hat Enterprise Virtualization Manager.

These parameters can be combined for specific commands. For example, **engine-log-collector --user=admin@internal --cluster ClusterA,ClusterB --hosts "SalesHost"*** specifies the user as **admin@internal** and limits the log collection to only **SalesHost** hosts in clusters **A** and **B**.

--no-hypervisors

Omits virtualization hosts from the log collection.

-u USER*, *--user=USER

Sets the user name for login. The *USER* is specified in the format *user@domain*, where *user* is the user name and *domain* is the directory services domain in use. The user must exist in directory services and be known to the Red Hat Enterprise Virtualization Manager.

-r FQDN*, *--rhevm=FQDN

Sets the fully qualified domain name of the Red Hat Enterprise Virtualization Manager server from which to collect logs, where *FQDN* is replaced by the fully qualified domain name of the Manager. It is assumed that the log collector is being run on the same local host as the Red Hat Enterprise Virtualization Manager; the default value is **localhost**.

-c CLUSTER*, *--cluster=CLUSTER

Collects logs from the virtualization hosts in the nominated *CLUSTER* in addition to logs from the Red Hat Enterprise Virtualization Manager. The cluster(s) for inclusion must be specified in a comma-separated list of cluster names or match patterns.

-d DATACENTER*, *--data-center=DATACENTER

Collects logs from the virtualization hosts in the nominated *DATACENTER* in addition to logs from the Red Hat Enterprise Virtualization Manager. The data center(s) for inclusion must be specified in a comma-separated list of data center names or match patterns.

-H HOSTS_LIST*, *--hosts=HOSTS_LIST

Collects logs from the virtualization hosts in the nominated *HOSTS_LIST* in addition to logs from the Red Hat Enterprise Virtualization Manager. The hosts for inclusion must be specified in a comma-separated list of host names, fully qualified domain names, or IP addresses. Match patterns are also valid.

SOS Report Options

The log collector uses the JBoss SOS plugin. Use the following options to activate data collection from the JMX console.

--jboss-home=JBOSS_HOME

JBoss installation directory path. The default is **/var/lib/jbossas**.

--java-home=JAVA_HOME

Java installation directory path. The default is **/usr/lib/jvm/java**.

--jboss-profile=JBOSS_PROFILE

Displays a quoted and space-separated list of server profiles; limits log collection to specified profiles. The default is ***'rhevm-slimmed'***.

--enable-jmx

Enables the collection of run-time metrics from Red Hat Enterprise Virtualization's JBoss JMX interface.

--jboss-user=JBOSS_USER

User with permissions to invoke JBoss JMX. The default is ***admin***.

--jboss-logsize=LOG_SIZE

Maximum size in MB for the retrieved log files.

--jboss-stdjar=STATE

Sets collection of JAR statistics for JBoss standard JARs. Replace *STATE* with **on** or **off**. The default is **on**.

--jboss-servjar=STATE

Sets collection of JAR statistics from any server configuration directories. Replace *STATE* with **on** or **off**. The default is **on**.

--jboss-twiddle=STATE

Sets collection of twiddle data on or off. Twiddle is the JBoss tool used to collect data from the JMX invoker. Replace *STATE* with **on** or **off**. The default is **on**.

the JMX invoker. Replace *STATE* with **on** or **off**. The default is **on**.

--jboss-appxml=XML_LIST

Displays a quoted and space-separated list of applications with XML descriptions to be retrieved. Default is **all**.

SSH Configuration

--ssh-port=PORT

Sets *PORT* as the port to use for SSH connections with virtualization hosts.

-k KEYFILE*, *--key-file=KEYFILE

Sets *KEYFILE* as the public SSH key to be used for accessing the virtualization hosts.

--max-connections=MAX_CONNECTIONS

Sets *MAX_CONNECTIONS* as the maximum concurrent SSH connections for logs from virtualization hosts. The default is **10**.

PostgreSQL Database Options

The ***pg-pass*** parameter includes the Red Hat Enterprise Virtualization Manager database in the log. The database user name and database name must be specified if they have been changed from the default values.

Use the ***pg-dbhost*** parameter if the database is not on the local host. Use the optional ***pg-host-key*** parameter to collect remote logs. The PostgreSQL SOS plugin must be installed on the database server for remote log collection to be successful.

--no-postgresql

Disables collection of database. Database collection is performed by default.

--pg-user=USER

Sets *USER* as the user name to use for connections with the database server. The default is **postgres**.

--pg-dbname=DBNAME

Sets *DBNAME* as the database name to use for connections with the database server. The default is **rhevm**.

--pg-dbhost=DBHOST

Sets *DBHOST* as the host name for the database server. The default is **localhost**.

--pg-host-key=KEYFILE

Sets *KEYFILE* as the public identity file (private key) for the database server. This value is not set by default; it is required only where the database does not exist on the local host.

See Also:

- Section B.5.1, “Log Collector”
- Section B.5.3, “Basic Log Collector Usage”

Report a bug

B.5.3. Basic Log Collector Usage

When the **engine-log-collector** command is run without specifying any additional parameters, its default behavior is to collect all logs from the Red Hat Enterprise Virtualization Manager and its attached hosts. It will also collect database logs unless the ***--no-postgresql*** parameter is added. In the following example, log collector is run to collect all logs from the Red Hat Enterprise Virtualization Manager and three attached hosts.

Example B.9. Log Collector Usage

```
# engine-log-collector
INFO: Gathering oVirt Engine information...
INFO: Gathering PostgreSQL the oVirt Engine database and log files from
localhost...
Please provide REST API password for the admin@internal oVirt Engine
user (CTRL+D to abort):
About to collect information from 3 hypervisors. Continue? (Y/n):
INFO: Gathering information from selected hypervisors...
INFO: collecting information from 192.168.122.250
INFO: collecting information from 192.168.122.251
INFO: collecting information from 192.168.122.252
INFO: finished collecting information from 192.168.122.250
INFO: finished collecting information from 192.168.122.251
INFO: finished collecting information from 192.168.122.252
Creating compressed archive...
INFO Log files have been collected and placed in
/tmp/logcollector/sosreport-rhn-account-20110804121320-ce2a.tar.xz.
The MD5 for this file is 6d741b78925998caff29020df2b2ce2a and its size
is 26.7M
```

See Also:

- Section B.5.1, “Log Collector”
- Section B.5.2, “Syntax for engine-log-collector Command”

Report a bug

B.6. SPICE Proxy

B.6.1. SPICE Proxy Overview

The SPICE Proxy is a tool used to connect SPICE Clients to guests when the SPICE Clients are outside the network that connects the hypervisors.

Setting up a SPICE Proxy consists of installing **Squid** on a machine and configuring **iptables** to allow proxy traffic through the firewall.

Turning a SPICE Proxy on consists of using **engine-config** on the Manager to set the key **SpiceProxyDefault** to a value consisting of the name and port of the proxy.

Turning a SPICE Proxy off consists of using **engine-config** on the Manager to remove the value that the key **SpiceProxyDefault** has been set to.

Report a bug

B.6.2. SPICE Proxy Machine Setup

Summary

This procedure explains how to set up a machine as a SPICE Proxy. A SPICE Proxy makes it possible to connect to the Red Hat Enterprise Virtualization network from outside the network. We use **Squid** in this procedure to provide proxy services.

Procedure B.2. Installing Squid on a RHEL Machine

1. Install **Squid** on the Proxy machine:

```
# yum install squid
```

2. Open **/etc/squid/squid.conf**. Change

```
http_access deny CONNECT !SSL_ports
```

   To

```
http_access deny CONNECT !Safe_ports
```

3. Restart the proxy:

```
# service squid restart
```

4. Open the default squid port:

```
# iptables -A INPUT -p tcp --dport 3128 -j ACCEPT
```

5. Make this iptables rule persistent:

```
# iptables-save
```

Result

You have now set up a machine as a SPICE proxy. Before connecting to the Red Hat Enterprise Virtualization network from outside the network, activate the SPICE proxy.

Report a bug

B.6.3. Turning on SPICE Proxy

Summary

This procedure explains how to activate (or turn on) the SPICE proxy.

Procedure B.3. Activating SPICE Proxy

1. On the Manager, use the engine-config tool to set a proxy:

```
# engine-config -s SpiceProxyDefault=someProxy
```

2. Restart the **ovirt-engine** service:

```
# service ovirt-engine restart
```

The proxy must have this form:

```
protocol://[host]:[port]
```

Note

Only the http protocol is supported by SPICE clients. If https is specified, the client will ignore the proxy setting and attempt a direct connection to the hypervisor.

Result

SPICE Proxy is now activated (turned on). It is now possible to connect to the Red Hat Enterprise Virtualization network through the SPICE proxy.

Report a bug

B.6.4. Turning Off a SPICE Proxy

Summary

This procedure explains how to turn off (deactivate) a SPICE proxy.

Procedure B.4. Turning Off a SPICE Proxy

1. Log in to the Manager:

```
$ ssh root@[IP of Manager]
```

2. Run the following command to clear the SPICE proxy:

```
# engine-config -s SpiceProxyDefault=""
```

3. Restart the Manager:

```
# service ovirt-engine restart
```

Result

SPICE proxy is now deactivated (turned off). It is no longer possible to connect to the Red Hat Enterprise Virtualization network through the SPICE proxy.

Report a bug

B.7. Squid Proxy

B.7.1. Installing and Configuring a Squid Proxy

Summary

This section explains how to install and configure a Squid Proxy to the User Portal.

Procedure B.5. Configuring a Squid Proxy

1. **Obtaining a Keypair**

   Obtain a keypair and certificate for the HTTPS port of the Squid proxy server.

   You can obtain this keypair the same way that you would obtain a keypair for another SSL/TLS service. The keypair is in the form of two PEM files which contain the private key and the signed certificate. In this document we assume that they are named **`proxy.key`** and **`proxy.cer`**.

   The keypair and certificate can also be generated using the certificate authority of the oVirt engine. If you already have the private key and certificate for the proxy and do not want to generate it with the oVirt engine certificate authority, skip to the next step.

2. **Generating a Keypair**

   Decide on a host name for the proxy. In this procedure, the proxy is called **`proxy.example.com`**.

   Decide on the rest of the distinguished name of the certificate for the proxy. The important part here is the "common name", which contains the host name of the proxy. Users' browsers use the common name to validate the connection. It is good practice to use the same country and same organization name used by the oVirt engine itself. Find this information by logging in to the oVirt engine machine and running the following command:

   ```
   [root@engine ~]# openssl x509 -in /etc/pki/ovirt-engine/ca.pem -
   noout -subject
   ```

   This command will output something like this:

   ```
   subject= /C=US/O=Example Inc./CN=engine.example.com.81108
   ```

   The relevant part here is **`/C=us/O=Example Inc.`**. Use this to build the complete distinguished name for the certificate for the proxy:

   ```
   /C=US/O=Example Inc./CN=proxy.example.com
   ```

   Log in to the proxy machine and generate a certificate signing request:

   ```
   [root@proxy ~]# openssl req -newkey rsa:2048 -subj '/C=US/O=Example
   Inc./CN=proxy.example.com' -nodes -keyout proxy.key -out proxy.req
   ```

Note

The quotes around the distinguished name for the certificate are very important. Do not leave them out.

The command will generate the key pair. It is very important that the private key isn't encrypted (that is the effect of the -nodes option) because otherwise you would need to type the password to start the proxy server.

The output of the command looks like this:

```
Generating a 2048 bit RSA private key
.....................................................+++
.....................................................................
..............+++
writing new private key to 'proxy.key'
-----
```

The command will generate two files: **proxy.key** and **proxy.req**. **proxy.key** is the private key. Keep this file safe. **proxy.req** is the certificate signing request. **proxy.req** doesn't require any special protection.

To generate the signed certificate, copy the **private.csr** file to the oVirt engine machine, using the **scp** command:

```
[root@proxy ~]# scp proxy.req engine.example.com:/etc/pki/ovirt-
engine/requests/.
```

Log in to the oVirt engine machine and run the following command to sign the certificate:

```
[root@engine ~]# /usr/share/ovirt-engine/bin/pki-enroll-request.sh
--name=proxy --days=3650 --subject='/C=US/O=Example
Inc./CN=proxy.example.com'
```

This will sign the certificate and make it valid for 10 years (3650 days). Set the certificate to expire earlier, if you prefer.

The output of the command looks like this:

```
Using configuration from openssl.conf
Check that the request matches the signature
Signature ok
The Subject's Distinguished Name is as follows
countryName           :PRINTABLE:'US'
organizationName      :PRINTABLE:'Example Inc.'
commonName            :PRINTABLE:'proxy.example.com'
Certificate is to be certified until Jul 10 10:05:24 2023 GMT (3650
days)

Write out database with 1 new entries
Data Base Updated
```

The generated certificate file is available in the directory **/etc/pki/ovirt-engine/certs** and should be named **proxy.cer**. Copy this file to the proxy machine:

```
[root@proxy ~]# scp engine.example.com:/etc/pki/ovirt-
engine/certs/proxy.cer .
```

Make sure that both the **proxy.key** and **proxy.cer** files are present on the proxy machine:

```
[root@proxy ~]# ls -l proxy.key proxy.cer
```

The output of this command will look like this:

```
-rw-r--r--. 1 root root 4902 Jul 12 12:11 proxy.cer
-rw-r--r--. 1 root root 1834 Jul 12 11:58 proxy.key
```

You are now ready to install and configure the proxy server.

3. **Install the Squid proxy server package**

   Install this system as follows:

   ```
   [root@proxy ~]# yum -y install squid
   ```

4. **Configure the Squid proxy server**

   Move the private key and signed certificate to a place where the proxy can access them, for example to the **/etc/squid** directory:

   ```
   [root@proxy ~]# cp proxy.key proxy.cer /etc/squid/.
   ```

   Set permissions so that the "squid" user can read these files:

   ```
   [root@proxy ~]# chgrp squid /etc/squid/proxy.*
   [root@proxy ~]# chmod 640 /etc/squid/proxy.*
   ```

   The Squid proxy will connect to the oVirt engine web server using the SSL protocol, and must verify the certificate used by the engine. Copy the certificate of the CA that signed the certificate of the oVirt engine web server to a place where the proxy can access it, for example **/etc/squid**. The default CA certificate is located in the **/etc/pki/ovirt-engine/ca.pem** file in the oVirt engine machine. Copy it with the following command:

   ```
   [root@proxy ~]# scp engine.example.com:/etc/pki/ovirt-engine/ca.pem
   /etc/squid/.
   ```

   Make sure that the "squid" user can read that file:

   ```
   [root@proxy ~]# chgrp squid /etc/squid/ca.pem
   [root@proxy ~]# chmod 640 /etc/squid/ca.pem
   ```

   If SELinux is in enforcing mode, change the context of port 443 using the **semanage** tool. This permits Squid to use port 443.

   ```
   [root@proxy ~]# yum install -y policycoreutils-python
   ```

```
[root@proxy ~]# semanage port -m -p tcp -t http_cache_port_t 443
```

Replace the existing squid configuration file with the following:

```
https_port 443 key=/etc/squid/proxy.key cert=/etc/squid/proxy.cer
ssl-bump defaultsite=engine.example.com
cache_peer engine.example.com parent 443 0 no-query originserver
ssl sslcafile=/etc/squid/ca.pem name=engine
cache_peer_access engine allow all
ssl_bump allow all
http_access allow all
```

5. **Restart the Squid Proxy Server**

   Run the following command in the proxy machine:

   ```
   [root@proxy ~]# service squid restart
   ```

6. **Configure the websockets proxy**

   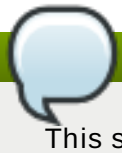

   Note

   This step is optional. Do this step only if you want to use the noVNC console or the Spice HTML 5 console.

   To use the noVNC or Spice HTML 5 consoles to connect to the console of virtual machines, the websocket proxy server must be configured on the machine on which the engine is installed. If you selected to configure the websocket proxy server when prompted during installing or upgrading the engine with the **`engine-setup`** command, the websocket proxy server will already be configured. If you did not select to configure the websocket proxy server at this time, you can configure it later by running the **`engine-setup`** command with the following option:

   ```
   engine-setup --otopi-
   environment="OVESETUP_CONFIG/websocketProxyConfig=bool:True"
   ```

   You must also make sure that the **`ovirt-websocket-proxy`** service is started and will start automatically on boot:

   ```
   [root@engine ~]# service ovirt-websocket-proxy status
   [root@engine ~]# chkconfig ovirt-websocket-proxy on
   ```

   Both the noVNC and the Spice HTML 5 consoles use the websocket protocol to connect to the virtual machines, but squid proxy server does not support the websockets protocol, so this communication cannot be proxied with Squid. Tell the system to connect directly to the websockets proxy running in the machine where the engine is running. To do this, update the **`WebSocketProxy`** configuration parameter using the "engine-config" tool:

   ```
   [root@engine ~]# engine-config \
   -s WebSocketProxy=engine.example.com:6100
   [root@engine ~]# service ovirt-engine restart
   ```

Important

If you skip this step the clients will assume that the websockets proxy is running in the proxy machine, and thus will fail to connect.

7. **Connect to the user portal using the complete URL**

   Connect to the User Portal using the complete URL, for instance:

```
https://proxy.example.com/UserPortal/org.ovirt.engine.ui.userportal
.UserPortal/UserPortal.html
```

Note

Shorter URLs, for example **`https://proxy.example.com/UserPortal`**, will not work. These shorter URLs are redirected to the long URL by the application server, using the 302 response code and the Location header. The version of **Squid** in Red Hat Enterprise Linux and Fedora (**Squid** version 3.1) does not support rewriting these headers.

Summary

You have installed and configured a Squid proxy to the User Portal.

Report a bug

Revision History

Revision 3.3-44 **Fri 20 Mar 2015** **Tahlia Richardson**

BZ#1203488 - Removed references to now-defunct Hypervisor Deployment Guide.

Revision 3.3-43 **Tue 07 Oct 2014** **Julie Wu**

BZ#1145040 - Added a note referencing the RHEL Security Guide.

Revision 3.3-42 **Thu 22 May 2014** **Andrew Burden**

Incorrect admonition regarding Self-Hosted Engine removed.

Revision 3.3-41 **Wed 5 Mar 2014** **Andrew Dahms**

BZ#1073160 - Updated the file names for the VirtIO virtual diskettes.

Revision 3.3-40 **Tue 4 Mar 2014** **Andrew Burden**

BZ#1072491 - Added a topic outlining the limitations of using a self-hosted engine.
BZ#978108 - Updated Self-Hosted Engine procedures.

Revision 3.3-39 **Mon 3 Mar 2014** **Andrew Dahms**

BZ#1071701 - Added a note to the topic on installing the self-hosted engine to enable the channel that provides the vdsm package.

Revision 3.3-38 **Mon 17 Feb 2014** **Andrew Dahms**

BZ#1059772 - Updated old references to beta channels and packages with full release versions.
BZ#973962 - Added a step for enabling the RHEV Management Agents repository for Red Hat Enterprise Linux hosts.

Revision 3.3-37 **Fri 14 Feb 2014** **Andrew Dahms**

BZ#1064137 - Updated the procedure for manually configuring a PostgreSQL database.
BZ#1059522 - Added a path to installation media in the procedure for configuring the self-hosted engine.
BZ#1047672 - Added a note that the VDSM service must be restarted after libvirt is updated.

Revision 3.3-36 **Fri 07 Feb 2014** **Andrew Dahms**

BZ#1063498 - Corrected the command for disabling repos in the procedure for upgrading to versio 3.3.
BZ#1062795 - Updated the method for configuring a websocket proxy server.
BZ#1062793 - Corrected the order of the procedure for upgrading to version 3.3.
BZ#1062791 - Generalized the final step in the procedure for upgrading to version 3.3.
BZ#1062719 - Added a note that users installing a self-hosted engine setup must follow the instructions in the relevant section.
BZ#1061761 - Changed 'authorized_hosts' to 'authorized_keys' in the procedure for adding Linux hosts.
BZ#1058934 - Corrected the RHN channel to disable during the upgrade from Red Hat Enterprise Virtualization 3.2 to 3.3.
BZ#1058353 - Corrected the packages for installing the Manager on 'my-engine' when configuring self-hosted engine.
BZ#1057366 - Changed a reference to 'hypervisor' in the introduction to the self-hosted engine to 'Red Hat Enterprise Virtualization'.
BZ#1057181 - Updated the channel and package required for installing the livecd-iso-to-disk tool.
BZ#1053101 - Changed references to Red Hat Enterprise Linux version 6 to version 6.5.
BZ#1039968 - Changed the CIDR notation for postgresql database configuration to '32'.

Revision 3.3-35 **Thu 23 Jan 2014** **Andrew Dahms**

BZ#1059506 - Added two steps to the procedure for manually configuring a postgresql database.
BZ#1055863 - Updated the syntax for creating and using answer files in conjunction with engine-setup.

Revision 3.3-34 **Thu 09 Jan 2014** **Andrew Dahms**

BZ#973962 - Added details on enabling the RHEV Management Agents repo in subscription manager.

Revision 3.3-33 **Tue 07 Jan 2014** **Zac Dover**

Publican 4 test

Revision 3.3-32 **Mon 23 Dec 2013** **Andrew Burden**

Updating Revision History for test build.

Revision 3.3-31 **Wed 18 Dec 2013** **Andrew Dahms**

BZ#1045269 - Updated the output from and description of the livecd-iso-to-disk utility.
BZ#1044876 - Updated the procedure for and output from configuring Hypervisors.
BZ#1044852 - Updated the procedure for and output from installing Hypervisors from boot media.
BZ#1044340 - Edited the process for upgrading to RHEVM 3.3 for clarity.
BZ#1044289 - Updated the standard output for and syntax of the engine-log-collector command.
BZ#1044251 - Updated the standard output for the engine-iso-uploader command.
BZ#1036425 - Revised the section on supported protocols in SPICE proxy configuration.
BZ#978198 - Generating and specifying passwords in Red Hat Enterprise Virtualization.

Revision 3.3-30 **Tue 17 Dec 2013** **Andrew Dahms**

BZ#1039450 - Updated the list of supported bonding modes for network interface cards.
BZ#976566 - Updated the New Host and Edit Host settings to include details on Foreman host provider hosts.
BZ#977619 - Support for Gluster processes on RHS nodes.

Revision 3.3-29 **Tue 17 Dec 2013** **Zac Dover**

BZ#962529 - RHEVM + SAM causes a conflict that causes #yum upgrade rhevm to fail - troubleshooting section added.

Revision 3.3-28 **Mon 16 Dec 2013** **Andrew Dahms**

BZ#969575 - Added a note outlining limitations on shutting down RHEL 3 and RHEL 4 systems.

Revision 3.3-27 **Fri 13 Dec 2013** **Andrew Dahms**

BZ#1042694 - Updated the procedure for adding a new cluster.
BZ#1042676 - Changed the syntax of all command-line tools from 'rhevm-' to 'engine-'.
BZ#981063 - Removed warnings that the rhevm network cannot be edited after addition of a cluster

Revision 3.3-26 **Thu 12 Dec 2013** **Andrew Burden**

BZ#978108 - Added a new chapter, 'Self-Hosted Engine', with six new topics.

Revision 3.3-25 **Wed 11 Dec 2013** **Andrew Dahms**

BZ#1039795 - Updated the syntax, output and procedure for running 'engine-cleanup'.
BZ#1039792 - Updated the syntax, output and procedure for running 'engine-setup'.
BZ#1039768 - Updated the 'subscription-manager subscribe' command syntax to 'subscription-manager attach'.
BZ#1039273 - Changed all references to 'User Principle Name' to 'User Principal Name'.
BZ#1039232 - Removed reference of adding a JBoss entitlement.
BZ#1038725 - Added a line outlining the need to restart the ovirt-engine service on turning off a SPICE proxy.
BZ#1018399 - Added a line outlining the need to disable NetworkManager.
BZ#984747 - Added a note to use a base installation of RHEL for RHEVM installations.
BZ#975683 - Updated the process for and output of the procedure for installing reports under otopi
BZ#975680 - Updated the process for and output of the procedure for installing dwh under otopi.
BZ#974367 - Changed references to 'Red Hat Enterprise Manager' to 'Red Hat Enterprise Virtualization Manager'.

Revision 3.3-24 **Fri 06 Dec 2013** **Andrew Dahms**

BZ#1030004 - Changed references to 'version.number' to '3.3'.

Revision 3.3-23 **Mon 25 Nov 2013** **Zac Dover**

BZ#978641 - Squid (User Portal Proxy)

Revision 3.3-22 **Mon 25 Nov 2013** **Andrew Dahms**

BZ#975668 - Outlined the use of otopi in installation and updated screen output.
BZ#975639 - Updated procedure for upgrading to version 3.3.

Revision 3.3-21 **Fri 22 Nov 2013** **Andrew Dahms**

BZ#978720 - Added procedure for importing external provider networks and an outline of the limitations thereof.

Revision 3.3-20 **Mon 18 Nov 2013** **Andrew Dahms**

BZ#960581 - Added a step to the 3.3 upgrade process to outline the procedure for updating EAP to version 6.1.

Revision 3.3-19 **Wed 13 Nov 2013** **Red Hat Engineering Content Services**

Completed line edits on the Installation Guide.

Revision 3.3-18 **Wed 13 Nov 2013** **Jodi Biddle**

BZ#978677 - Added 3.3 upgrade topic.

Revision 3.3-17 **Thu 17 Oct 2013** **Red Hat Engineering Content Services**

BZ#719765 - Added processing rule to Setting up a Virtualization Host Logging Server. Now, local messages are processed separately from received messages.

Revision 3.3-16 **Tue 15 Oct 2013** **Zac Dover**

BZ#978157 - Support ISO domain on any file storage

Revision 3.3-15 **Fri 11 Oct 2013** **Red Hat Engineering Content Services**

BZ#975620 - updated Installing and Configuring the History Database to reflect automatic creation of read-only user for safe remote access.
BZ#978708 - Added GlusterFS storage domains to the Introduction to Storage content.

Revision 3.3-14	Wed 09 Oct 2013	Zac Dover

BZ#978875 - pNFS Client Support - bumping the revnumber so that the book travels through the buildsystem.

Revision 3.3-13	Fri 04 Oct 2013	Red Hat Engineering Content Services

BZ#978220- Fixed typos in bonding logic topic.
BZ#976187 - Updated the title of "Designate a Specific Traffic Type for a Logical Network with the Manage Networks Window".
BZ#976187 - Added migration network as a network type.
Replaced screenshot in "Designate a Specific Traffic Type" topic.
BZ#975298 - Updated Host General Settings topic to include SSH public key and SSH fingerprint.
Updated "Configuring Hypervisor Management Server" topic, changed "Verfify" to "Retrieve".
BZ#975620 - Added note to rhevm-dwh-setup steps about opening TCP port 5432.
BZ#975620 - Added rhevm-dwh-setup steps that configure read-only access to history database.
BZ#994235 - Updated the "Changing the Cluster Compatibility Version" topic to make it clearer that hosts must be upgraded first, and that maintenance mode is not required.

Revision 3.3-12	Wed 02 Oct 2013	Andrew Burden

Updated 'Edit Management Network' window title and content; associated with BZ#978092

Revision 3.3-11	Mon 30 Sep 2013	Red Hat Engineering Content Services

BZ#976251 - Finished table of features requiring a compatibility upgrade to Red Hat Enterprise Virtualization 3.3

Revision 3.3-10	Thu 26 Sep 2013	Charles Boyle

BZ#978714 - Add information about new engine log.
BZ#978854 - Changed JBoss Enterprise Application Platform to Red Hat JBoss Enterprise Application Platform.

Revision 3.3-9	Thu 29 Aug 2013	Red Hat Engineering Content Services

Added basic database configuration steps.
Adding information about manual database configuration.
Rewrote the "Removing the Red Hat Enterprise Virtualization Manager" topic in light of the otopi installation framework.
Updated URL used to access the Manager in "Connecting to the Administration Portal" topic.
Rewrote "Configuring Red Hat Enterprise Virtualization Manager" for 3.3 in light of move to otopi installer.

Revision 3.3-8	Fri 23 Aug 2013	Red Hat Engineering Content Services

BZ#978875 - pNFS support
BZ#978152 - Support for OpenLDAP as Directory Service
BZ#995846 - Removing an incorrect "presently"

Revision 3.3-7	Thu 22 Aug 2013	Red Hat Engineering Content Services

Changed engine-config to rhevm-config in SPICE Proxy content.

Revision 3.3-6	Fri 16 Aug 2013	Red Hat Engineering Content Services

BZ#995916 - SPICE Proxy documentation added

Revision 3.3-5	Tue 13 Aug 2013	Red Hat Engineering Content Services

BZ#988172 - Made it clear which SPICE features are supported by which clients.
BZ#986735 - Updated host firewall ports to reflect change SPICE range.

Revision 3.3-4	Mon 12 Aug 2013	Red Hat Engineering Content Services

BZ#981764 - Internet Explorer 10 touchscreen version not supported.
BZ#988172 - Updated supported SPICE features in Red Hat Enterprise Linux 6 clients.
BZ#719765 - Updated semanage command.
Updated Host Compatibility Matrix for 3.3 release.

Revision 3.3-3	Fri 09 Aug 2013	Red Hat Engineering Content Services

BZ#719765 - Updated selector to send specific messages to server.
BZ#719765 - Updated semanage command.
BZ#965332 - Changed suggested server type to install on Linux host from 6.3 to 6.
BZ#988172 - Note the limitations in SPICE feature support in RHEL 5.X
BZ#981764 - Added Windows 8 to list of SPICE clients.
Removed mention of "four supported bonding modes" as all modes are supported using custom bonds if not from the drop down menu.
Changed 'rhev-hypervisor' to 'rhev-hypervisor6'.
Edits to Win8/2012 SPICE support statement

Revision 3.3-2	Thu 01 Aug 2013	Red Hat Engineering Content Services

BZ#989346 - change to support matrix to highlight that SPICE is not supported on Win 8/2012
BZ#987207 - rhev-hypervisor corrected to rhev-hypervisor6
thildred: rhev 3.3 installation guide revision history

Revision 3.3-1	Thu 18 Jul 2013	Tim Hildred

Initial creation for 3.3 release.